

# Introduction to Industrial Organization

## Teaching Notes

Luís Cabral

### 1 What Is Industrial Organization?

As the title suggests, the purpose of this chapter is to introduce students to the concept of industrial organization. There are two main parts to this. First, to explain what the term "industrial" means. Much too often, readers assume that industrial organization (or industrial economics) is the study of the industrial sector, as opposed to agriculture or services. In fact, "industry" refers to market or activity sector, as in the insurance industry or the cement industry. A more appropriate name for IO, one that would avoid this confusion, might be "the economics of firms and markets" or "the economics of imperfect competition."

The second step in explaining IO consists of outlining the main issues addressed in IO. There are certainly many ways of organizing and classifying the main issues. The solution adopted in this chapter has the advantage of highlighting:

1. The double perspective of positive and normative analysis. Positive analysis (how things are) is reflected in the first question (is there market power) and the third one (the implications of market power); normative analysis (how things should be) is reflected in the second question (how firms maintain market power) and the fourth one (what policy makers can do about market power).
2. The double perspective of the firm and policy makers. The firm's perspective is reflected in the second question (how to sustain market power), whereas the policy maker's is reflected in the fourth question (what to do about market power).

The relative weight given to each question naturally depends on the particular audience of the book. For example, when teaching MBA students the emphasis should be put on the second question, whereas the fourth question is a central one when teaching economics majors.

The example chosen to illustrate some of the concepts in IO (Zantac) is one among many other possible examples. Alternatives include airlines, telecommunications, video games and artificial sweeteners, to name a few (all these are introduced at later stages in the book). Instructors may want to choose an example that they are familiar with. Whichever it is, the important thing is to start off with an example, so that students get the idea that the concepts they are being introduced to are useful to understand particular aspects of the reality of firms and markets.

## 2 Basic Microeconomics

This chapter is addressed to readers who have not previously studied microeconomics (or have done it so long ago that they need a refresher). When teaching economics majors, who will normally have taken at least one economics class before taking IO, this chapter may be skipped without much loss.

Short as it is, this chapter encapsulates some of the most important concepts in microeconomics. The concepts of consumer surplus, producer surplus, and allocative efficiency play a central role in any discussion regarding public policy. For this reason, in courses that emphasize the public policy side of IO, it is well worth it to spend some extra time making sure students understand these concepts. It is also important to emphasize that allocative efficiency is only part of the story. Historically, economists have placed a little too much weight on static allocative efficiency, leaving productive efficiency and dynamic efficiency to a secondary level.

In courses that emphasize the strategy aspect of IO (e.g., MBA courses), special emphasis should be put on the concept of economic cost and its applications for decision making. Rational decisions should take into account opportunity costs, even if these do not correspond to actual expenditures; and should exclude sunk costs, even if these correspond to actual expenditures. Although, in most of the book, this does not appear explicitly, the concept of economic cost is implicit in just about every economic decision.

In courses that include both the managerial and the strategy aspects of microeconomics (e.g., the MBA micro course at LBS and Berkeley), it might be worth to include additional materials to discuss the concept of economic costs: these ideas are much too important to be limited to a short chapter as this.

## 3 The Firm

As a first-order approximation, one may say that IO focuses primarily on competition between firms, placing relatively little weight on the internal workings of each firm. In particular, from the point of view of strategy, IO suggests how firms may improve their situation as a result of actions with respect to rivals, not so much actions with respect to its internal organization.

This is not a bad thing in and of itself: the effectiveness of every partial science results precisely from the fact that they are partial and are thus able to tackle a particular issue in depth. However, this is a limitation that should be duly acknowledged: there is more to the world of firms and markets than IO addresses.

This chapter highlights three aspects in which the IO view of the firm has to be taken with caution. First, IO assumes that firms behave as profit maximizers. Second, IO takes the limits of the firm—in fact, the very existence of firms—as exogenously given. Third, IO has little to say about why firms are different—other than to assume it. This chapter is therefore included in the introductory part of the book not because its contents are used later on. Rather, the idea is to help the student put the later chapters in their proper context.

The statement at the beginning of the first paragraph is perhaps a little too strong. Many IO economists have recently devoted themselves to studying the theory of the firm, one of the fastest growing areas of microeconomics. In fact, several schools now offer courses on the theory of the firm. When this is the case and students have taken or will take a course on the theory of the firm, the discussion of this chapter may be reduced or omitted.

Moreover, it is increasingly common for IO courses (or microeconomics courses with an emphasis on IO) to delve into issues of the theory of the firm, in particular the vertical boundaries of the firm. In that case, it may be worth to join together the discussion of Section 3.2 (what determines the size of the firm?) and that of Chapter 11 (vertical relations).

#### 4 Games and Strategy

One of the great accomplishments of IO has been the rigorous application of game theory to the analysis of market competition. In fact, the game theoretic nature of IO is one of the main reasons why IO became so popular within economics—and with other related disciplines, such as strategy and marketing.

This chapter, which concludes the introductory part of the book, presents the main concepts of game theory that are necessary to understand the rest of the text. In fact, the chapter goes a little beyond what is strictly required to follow the rest of the text. As indicated in footnote c, Sections 4.2 and 4.3 can be skipped in a first reading. Not that sequential games and repeated games are unimportant: in fact, Chapters 12 and 15 introduce various sequential games; and Chapter 8 is essentially based on the notion of repeated games. However, Chapters 8, 12 and 15 are written in a way that allows for its comprehension without the prior reading of Sections 4.2 and 4.3.

There are two possible “strategies” regarding the placement of this chapter. One is to present the material upfront in its entirety. This has the advantage of giving a more unified view of game theory; it has the disadvantage that students have to bear a fair amount of material without seeing a significant payoff in terms of applications. An alternative strategy is to present Section 4.1 together with Chapter 7, Section 4.2 together with Chapters 12 and 15, and Section 4.3 together with Chapter 8. This has the advantage of introducing each concept together with an application; it has the disadvantage that game theory is introduced in a somewhat fragmented way.

#### 5 Monopoly and Regulation

One of the important points this chapter attempts to make is that monopoly is not a zero-one concept: it all depends on how the market is defined. In fact, what really matters is not so much whether a firm dominates one hundred percent or less than one hundred percent of the market; the important thing is whether the firm has the power to set prices above marginal cost.

This distinction plays an important role in the way the whole book is structured. Many IO texts—e.g., Tirole’s *The Theory of Industrial Organization*—are divided into two parts: monopoly and oligopoly. Such distinction is not made here. Each issue (price discrimination, advertising, etc) is approached both from the monopoly and the oligopoly point of view. The underlying idea is that the important distinction is not so much between monopoly and oligopoly but rather between firms with market power and firms without market power. (This is not exactly right, but is probably a good first-order approximation.)

The above distinction is also important when it comes to discuss the issue of regulation. Regulation is called for not because a firm is a monopolist but rather because a firm commands

excessive market power—and even then the question has to be asked whether regulation does more good than bad (cf the Microsoft case).

It should finally be noted that the discussion of regulation in this chapter is very summary. In courses where regulation is given a bigger weight, instructors may want to assign additional readings on the subject. There are some excellent introductory texts covering regulation, for example, Harrington, Vernon and Viscusi, *The Economics of Regulation*, MIT Press.

## 6 Perfect (and Almost Perfect) Competition

One may think of the spectrum of possible market structures as a segment with two extremes: monopoly and perfect competition. Both of these extremes are ideal situations that do not exist *exactly* in reality. In particular, there is no real-world industry that satisfies all of the assumptions of the perfect competition model. How useful is the model, then?

To address this question, this chapter presents two models that relax some of the rather extreme assumptions of the perfect competition model: the models of competitive selection and monopolistic competition. Of these two models, the first one is a bit of a novelty in IO textbooks. The reason for its introduction is fourfold. First, the theoretical developments around this model have been one of the important developments in the field during the 1980s and 90s, though this fact is not always duly appreciated. Second, the advent of powerful computers and rich datasets has created a wealth of well documented stylized facts regarding the process of firm entry, growth and exit, among others. Third, the theoretical models seem to fit the facts fairly well (actually, the models are a great help in understanding the facts). Fourth, there are still a large number of theoretical points to be studied and facts to be investigated; that is, this is important area in terms of future research.

This chapter, together with the previous one (Monopoly and regulation), is supposed to set the stage for the rest of the book. Together, these chapters make the point that the extreme models are a good approximation to situations that are close to the ends of the spectrum. So, a dominant firm (close to a monopolist, but not quite) behaves in a way similar to a monopolist; a firm in a competitive selection or monopolistic competition context (close to perfect competition, but not quite) behaves in a way similar to a firm under perfect competition. This leaves a host of other possible intermediate market structures, those with more than one firm but less than the "large number" assumed by perfect competition. This is where oligopoly analysis enters: Chapter 7 and the following chapters.

## 7 Oligopoly Competition

Many IO textbooks introduce the theory of oligopoly by presenting a series of models with little explanation of what their purpose is. When this happens, instructors are often faced with questions such as: Why so many models? Why are we not given the "right" model only, saving us the trouble of going through all the other ones?

The proliferation of IO models started in the 19th Century when, responding to the early work by Cournot, Bertrand proposed an alternative model of oligopoly competition that implied strikingly different predictions from the early model. As a result, the debate about which model is "best"—Cournot or Bertrand—has evolved into a cottage industry within IO.

The approach taken in this chapter is one that reflects recent developments in IO theory. In particular, it views the Bertrand and Cournot models as particular cases of a more general model of oligopoly competition where firms choose prices *and* quantities (or capacities). Seeing things from this perspective, the Bertrand model arises as the "right" solution when firms can adjust capacities faster than prices (e.g., software), whereas the Cournot model corresponds to the opposite case, the case when prices can vary faster than capacities (e.g., wheat, cement).

This way of looking at things is important for this chapter and, more generally, for the whole of oligopoly theory. There is no single model that is better than all other models. Each model is a better description of a certain type of industry. The art (and science) is then to determine what is the appropriate model for each situation.

Another important point that this chapter attempts to put across is that models are useful, operational tools to analyze a particular real-world situation—oligopoly competition, in the case of IO. For this reason, a fairly long section on comparative statics is included at the end of the chapter. Comparative statics is one of the main goals of economic modeling. It allows one to estimate the impact of an exogenous change in the equilibrium values of a given system—in the case of IO, prices, quantities, market shares, and so on.

Even when "theoretical" lectures are complemented by "applied" sections, it is recommended that some time is spent in the main lectures to go through the material in Section 7.5. Among other things, this will have the benefit of making students more confident in the value of the models—and theory, more generally. As Chesterton so aptly put it, the best help for practical life is a good theory.

The Bertrand model implies a very striking result: even if there are only two competitors, prices will be set at the level of marginal cost. From the perspective of positive economics, this implies what we might call the "Bertrand paradox" (Tirole, 1989): in reality, there are many industries that look like the Bertrand model but where prices are higher—or much higher—than marginal cost. Chapter 7 includes a brief description of the three main explanations for this apparent paradox: capacity constraints (Chapter 7), dynamic interaction (Chapter 8) and product differentiation (Chapter 12).

From the perspective of competitive strategy, the outcome of the Bertrand model has aptly been described as the "Bertrand trap" (Hermalin, 1993): an outcome that firms should avoid at all costs. The three ways of explaining the Bertrand paradox (positive analysis) can then also be interpreted as ways of exiting the Bertrand trap (normative analysis)—especially the "escapes" described in Chapters 8 and 12. Depending on the nature of the audience, one or the other view may be emphasized. For example, when teaching MBA students it may make more sense to stress the idea of the Bertrand trap.

## 8 Collusion

As mentioned in Chapter 7, there are essentially three "solutions" to the Bertrand "paradox"—the fact that, under the assumptions of the Bertrand model, prices are equal to marginal cost even with only two competitors. Chapter 7 introduces the first of these solutions: capacity constraints. Chapters 8 and 12 focus on the two other solutions, repetition and product differentiation.

The main point of Chapter 8 is that, when firms compete over a number of periods, equilibrium prices may be above marginal cost when otherwise (i.e., with no repetition) the only

equilibrium would be to price at the level of marginal cost. The intuition is that the short-term gains from undercutting a competitor may not compensate the longer-term losses from entering into a price war.

In addition to explaining this idea, Chapter 8 outlines the main conditions under which we would expect pricing above marginal cost to be an equilibrium. The first set of conditions is that the discount factor is sufficiently close to one, which essentially means that firms interact frequently enough. Sections 8.2 and 8.3 address other factors that facilitate collusion.

Implicit or explicit collusion is an area of great importance in antitrust and competition policy. This is a chapter where it should be easy to find many current examples. In addition to the ones in the book, some are presented in this website. Moreover, a possible assignment for this chapter could be to ask students to come up with a current example featuring collusion and possible antitrust remedies. For this purpose, it may be useful to follow the links to "government institutions" under the links area of this website.

## 9 Market Structure and Market Power

This is the only chapter in the book that is primarily devoted to empirical work. It is a more research-oriented chapter than the remaining ones. For this reason, it can be omitted from courses that are relatively more strategy-oriented.

As in rest of the book, the focus is primarily on issues, rather than the literature. It should be noted, however, that there is a very extensive literature on the structure-performance relationship and on the empirical estimation of market power. Instructors interested in going deeper in these issues may want to consult the chapters by Schmalensee and Bresnahan in the Handbook of Industrial Organization. Martin's "Advanced Industrial Organization" also contains a good discussion of the empirical literature.

One possible assignment in courses that are research-oriented—and when students have a good knowledge of econometrics—is to perform some basic econometric analysis based on existing datasets. Some links to existing datasets may be found in the links section of this website.

## 10 Price Discrimination

This is the first of a series of four chapters that depart from the basic models of oligopoly and look at price and non-price strategies. This fourth part of the book is particularly important in courses with an emphasis on the strategy and marketing aspects of IO (especially the marketing aspects).

Various IO texts are divided into "monopoly" and "oligopoly" parts, with each part then subdivided according to the different issues. In this text, the classification is by issue, with each issue including both the monopoly and the oligopoly cases. (In this chapter, for the most part, we consider the monopoly case.)

There is a discussion regarding the classification of types of price discrimination, namely the classification, due to Pigou, of first-, second- and third-degree price discrimination. Although a paragraph is included in the text, it is probably best to avoid a discussion that is little more than semantics. Instead, the instructor can focus on the different instances of price

discrimination using more specific names, such as self-selection, group discrimination, and so on.

One possible assignment for this chapter is to provide students with a demand system and ask them to select prices to maximize profit; and then compare the solution chosen by each student with the optimal solution. Based on the author's experience, such exercise has a number of advantages. First, it gives students the opportunity to take a hands-on approach to the problem of price discrimination. Second, it greatly motivates students to learn what the optimal solution is. [I am hoping to post a demand system for this purpose some time soon.—LC]

## 11 Vertical Relations

Vertical relations are one of the areas of most intense research activity in recent years. Although few IO texts devote an entire chapter to this topic, the amount of work developed during the 1980s and 90s more than justifies its separate treatment.

From the point of view of public policy, the main issue in this chapter is the conflict between efficiency reasons and market power motivations for different forms of vertical relations, from simple vertical restraints to full-fledged vertical integration. Although the analysis of public policy towards vertical relations is left to the end of the chapter, it may be worth starting off with the presentation of this conflict as a way to motivate the chapter. This is especially true if the audience is primarily composed of economics majors who are interested in public policy issues. (The same consideration applies to the merger analysis, which is introduced in Section 15.3.)

Contract theory plays a central role in the study of vertical relations. In fact, many of the developments in the theory of vertical relations have appeared in parallel with developments in contract theory. Most of these developments are rather technical and were therefore excluded from the main text. A series of supplemental sections present some of the main results, as well as bibliographic references to more thorough presentations of the material.

As suggested above, vertical relations are closely related to the topic of vertical integration, which is discussed in Section 3.2. (In vertical integration, we can also find the dichotomy between efficiency considerations and market power considerations.) In fact, one possibility is to discuss both topics at the same time, under the heading vertical restraints and vertical integration, or something similar.

## 12 Product Differentiation

As mentioned in Chapter 7, product differentiation is one of the possible answers to the Bertrand "paradox," the fact that, even with two players only, price competition implies zero margins. As in many other chapters, one may approach this problem from a "positive" point of view or from a "normative" point of view. From a positive point of view, the question is how to solve the "Bertrand paradox": how come in practice we observe firms pricing above marginal cost even though theory seems to imply pricing at the level of marginal cost? From a normative point of view, the relevant issue is how firms can escape the "Bertrand trap," the trap of engaging in aggressive price-cutting. In strategy-oriented courses, the emphasis should be on the latter aspect. This is especially true with respect to the topic of product positioning.

In most IO texts, one of the central aspects in the chapter on product differentiation is the distinction between vertical and horizontal product differentiation. While this is a useful distinction for pedagogical reasons, in practice most cases combine both horizontal and vertical product differentiation. The general model that encompasses both forms of differentiation is the hedonic, or characteristics, model of product differentiation. Although the characteristics model is difficult to implement, the basic idea is quite simple. It is therefore worth it to take students through a simplified example—as the one in the text—so that they get an idea of what’s involved in applying the model.

The inclusion of imperfect information and switching costs in the same chapter as product differentiation is somewhat unusual. There are two reasons that justify this option. First, both product differentiation and imperfect information / switching costs share the feature that cross price elasticity is lower than it would be in the situation of product homogeneity / perfect information / no switching costs. Second, the actual models that describe the situations of imperfect information and switching costs are almost isomorphic (i.e., identical) to the model of product differentiation.

### 13 Advertising

Advertising is not one of the core topics of IO—at least, not of IO narrowly defined. This chapter may therefore be omitted from a course directed primarily at economics majors. However, as mentioned in the preface, IO has gradually been expanding to the point of overlapping significantly with fields like marketing and strategy, which justifies the inclusion of advertising as a separate topic.

One of the central issues in the IO analysis of advertising is the relation between advertising competition and price competition. Not much economics research has been done on this topic. As a consequence, most of the points presented in section 13.3 are rather tentative: more questions are raised than answered. (This is one of the sections of the text where the issue-driven, rather than literature-driven, nature of the text is most evident.)

### 14 Entry Costs, Market Structure, and Welfare

This chapter and the next close the “loop” implicit in the structure-conduct-performance paradigm. Whereas previous chapters take structure as exogenously determined, Chapters 14 and 15 look at the endogenous determination of market structure. Chapter 14 focuses on the role of entry costs. In particular, it introduces the distinction between exogenous and endogenous entry costs (a distinction that is absent from most other IO texts).

The public policy issue addressed in this chapter is that of the welfare effects of entry. An important strand of the economics literature on this topic has emphasized the fact that, in a world of second best, free entry may not be a good thing; in particular, free entry may lead to excessive entry. This literature overstates the importance of the business-stealing effect (the externality between entrants and incumbents) and underemphasizes the importance of firm heterogeneity, namely the fact free entry may help replacing inefficient firms with more efficient ones. To compensate for this perceived bias, Chapter 14 introduces the discussion of firm heterogeneity and free entry, including the example of deregulation and entry in the U.S. telecoms industry. (The importance of firm heterogeneity is also underappreciated in the

treatment of competitive markets, a fact that justifies the inclusion of Section 6.2. On telecoms deregulation, see also Chapter 5.)

## 15 Strategic Behavior, Entry and Exit

This is the longest chapter of the book. In fact, a case might be made for the division of this chapter into three different chapters, corresponding to each of the three sections.

In strategy-oriented courses, the emphasis should be placed on Sections 15.1 and 15.2. In policy-oriented courses, by contrast, the emphasis should be on Section 15.3.

This chapter differs from standard IO presentations in several respects:

1. In the treatment of entry deterrence, the emphasis is normally placed on capacity expansion as a means of entry deterrence. However, aside from titanium dioxide and a few other examples, this does not seem to be a very common form of entry deterrence. By contrast, product proliferation and contracts seem to be a far more prevalent strategy. Accordingly, Section 15.1 includes a treatment of the latter two forms of entry deterrence.
2. In the treatment of predation, the emphasis is normally placed on asymmetric information as a justification for rational predatory pricing. But the deep purse or the market share explanations for predation seem empirically more relevant. Accordingly, more emphasis is placed on these alternative theories than in standard IO presentations of the material.
3. Merger waves are empirically a very relevant phenomenon. Almost every week, one finds at least one instance of a merger that either follows or precedes other mergers / merger announcements in the same industry. This raises a number of important issues, both from a positive and from a normative point of view. Unfortunately, not much economics research has been done of these issues—and not much is said in standard IO textbook presentations. As in other parts of the book, Section 15.3 raises more questions than it answers.

## 16 Research and Development

From a positive point of view, the central question addressed in this chapter is the relation between market structure and R&D: first, how market structure affects the level and direction of R&D; second, how R&D determines the evolution of market structure. The first part of the central question has been given extensive treatment, both theoretically and empirically. (On the empirical literature, see Levin's chapter in the "Handbook of Industrial Organization" and Sutton's critique in his recent "Technology and Market Structure".) The second part of the question—the dynamics of R&D and market structure—is relatively more recent in terms of economics research. Specifically, one central question is whether the dynamics of R&D contribute to a process of increasing dominance (whereby market leaders increase their lead and monopolists persist as such) or something like what Schumpeter referred to as the process of creative destruction (whereby new firms constantly replace incumbents). There is an interesting parallel between this question and the macroeconomics literature on growth and convergence (do rich countries tend to become even richer or do poor countries tend to catch up?).

From a normative point of view, there are two important questions. First, the antitrust treatment of R&D agreements between firms. Public policy has normally been distrustful of agreements between firms. However, there are reasons to make R&D an exception to this general stance.

The second important normative question refers to the protection of intellectual property (IP). The basic trade-off here is between maximizing allocative efficiency and providing the right incentives for investment in R&D (in other words, the trade-off between static and dynamic efficiency). The development of the internet has created a host of novel IP issues, both related to patent law and to copyright law (should Amazon.com be given a patent for "one-click shopping"? Did Napster.com violate copyright law, and, more fundamentally, What should the copyright law state?). The economic analysis of these issues will certainly constitute an important part of the economics literature on R&D in the years to come.

## **17 Networks and Standards**

This is the most innovative chapter in the book (innovative with respect to standard IO presentations). The field of economics of networks and standards has developed over the past 15 years or so, and is only gradually making its way into textbooks. The number of industries with direct or indirect (a.k.a. virtual) network effects is extremely large. This is especially true among high-tech industries. The treatment of networks and standards in a separate chapter is therefore amply justified.

In strategy-oriented courses, a special emphasis should be put on the topic of strategic compatibility decisions. Although the economic analysis has still not settled on this issue, this is a very important issue from the point of view of strategy, in fact one where new examples keep constantly coming up (witness the recent DVD battles).

From the point of view of public policy, the fundamental issue is the role of the government in the process of standardization. The text refers to an important real-world case study: the contrast between the European approach and the U.S. approach to standardization in wireless telephony (second and third generation). This is a process that is still not over; it may therefore be interesting to follow the recent developments in it.