ONE

Introduction and Overview

We present and compare various techniques for estimating *market power* – the ability to set price profitably above marginal cost – and *strategies* – game-theoretic plans used by firms to compete with rivals. We start by examining static model approaches to estimating market power. Then, we extend our market power analysis to dynamic models. Finally, we develop methods to estimate firms' strategies directly and examine how these strategies determine market power. Throughout, we emphasize that the type of study we undertake depends crucially on which variables are observed and on the unobserved strategies of firms.

Research on market power and strategies is important for both policy makers and academics. It provides evidence that policy makers can use to improve antitrust and merger laws. It is used in court cases – and presumably will be increasingly employed in the future. It also allows academics to test theoretical models that were previously accepted on faith.

THREE MAIN QUESTIONS

Throughout this book, we focus on three major questions:

(1) How much market power does a firm or industry exercise?
(2) What are the major factors that determine this market power?
(3) How do firms' strategies determine market power?

Whether one can observe a reliable measure of market power often determines how researchers proceed. For example, in a one-period market, if we can observe the price, p, charged and can observe or estimate the marginal cost, MC, we can provide a direct answer to the first question – how much market power is exercised – by measuring the gap between price and marginal cost. To make this answer

independent of the units of measurement of price, p, and marginal cost, MC, it is traditional to use the Lerner's (1934) index,

$$L \equiv \frac{p - MC}{p}. \tag{1.1}$$

Lerner's index is the percentage markup of price over marginal cost. This construct answers the question "How much market power does a firm exercise?" and not the question "How much market power does a firm (in theory) possess?" In a competitive market, price equals marginal cost, so any gap – a positive Lerner's index – shows that firms are exercising market power.

As we discuss in Chapter 2, most *structure–conduct–performance* (SCP) studies implicitly treat Question 1 as moot by assuming that an available measure reflects the degree of market power in an industry.[1] SCP researchers focus on answering our second question by explaining variations in market power across industries based on the structure of the industry and other factors.

In Chapters 3 and 4, we discuss the static model approach that starts by rejecting the assumption of the SCP studies that we have reliable measures of Lerner's index or other measures of market power. These studies answer Question 1 – the degree of market power – and Question 2 – the cause of market power – simultaneously using formal models that reflect exogenous institutional features of a particular industry.

In dynamic models, even if we observe current price and some measure of short-run marginal cost, we generally cannot determine the degree of market power, which may depend on unobserved opportunity costs or option values. In Chapters 5 through 8, we examine the answers to Questions 1 and 2 using dynamic structural models. In these chapters, we discuss how the cause of the dynamics plays a critical role in determining the appropriate model. In many dynamic studies, the single-period models are generalized to allow firms to have multiperiod strategies.

Finally, in Chapters 9 and 10, we discuss methods to estimate firms' strategies directly, Question 3, along with demand and cost functions. We show how to use these estimated strategies to solve for the market equilibrium, which allows us to estimate market power.

We now provide a detailed overview of the book. The book is divided into four parts. We start by providing a basic background chapter on the traditional structure–conduct–performance (SCP) approach. We then show how to estimate static models, dynamic models, and strategies. The SCP, static, and dynamic analyses focus on our first two questions: measuring and explaining market power. In the final section, where we estimate firm strategies – answer our third question – we show how to use those estimated strategies to provide good answers to the first two questions as well. The end of the book contains a detailed statistical appendix, which provides a background on all the traditional estimation methods we use,

[1] In contrast to legal discussions, we use the terms "industry" and "market" interchangeably.

and a relatively complete presentation of information theoretic methods that we use in the strategy section.

STRUCTURE–CONDUCT–PERFORMANCE

The traditional approach to empirical studies of market power, Mason's (1939, 1949) SCP approach, has been in use for two-thirds of a century.[2] It holds that an industry's performance (the success of an industry in producing benefits for consumers or profits for firms) depends on the conduct (behavior) of sellers and buyers, which depends on the structure of the market. The structure is often summarized by the number of firms or some other measure of the distribution of firms, such as the relative market shares of the largest firms.

Mason and his colleagues initially conducted case studies of individual industries. Eventually, Bain (1951, 1956) and others introduced comparisons across industries. Since Bain's early work, a typical SCP study regresses a measure of market power such as profit or the gap between price and a cost measure on a structural variable such as the market shares of the four largest firms and other variables. That is, these studies presume that the answer to our first question – how much market power is exercised – is known, and concentrate on the second question – what causes this market power.

This literature has been criticized as being descriptive rather than theoretical. A more generous way to view it is that SCP researchers are estimating a reduced-form regression in which performance is linked to market structure. Although these researchers rarely use formal theoretical models to justify their empirical research, many such models exist.

For simplicity, suppose that the n firms in a market produce a homogeneous product, so that industry output, Q, equals sum of the output of the n firms: $Q = q_1 + q_2 + \cdots + q_n$. Given full information and no other distortions, there is a single market price, p, determined by the inverse market demand curve, $p(Q)$. The single-period profit of Firm i is

$$\pi_i = p(Q)q_i - mq_i,$$

where m is a constant marginal cost that is common to all firms.

The firms engage in a one-period game. Each firm has a strategy determining its actions. A set of strategies is a Nash equilibrium if, holding the strategies of all other players (firms) constant, no player (firm) can obtain a higher payoff (profit) by choosing a different strategy. In particular, suppose that the firms choose quantities. We call the outcome of this game a Nash-in-quantities, Cournot, or Nash-Cournot equilibrium.

[2] Schmalensee (1989) thoroughly analyzed the SCP literature through the 1980s.

Each firm's first-order condition, given the Nash-in-quantities assumption that other firms do not change their outputs, is obtained by setting the derivative of profit with respect to the firm's output equation to zero: $\mathrm{d}\pi_i/\mathrm{d}q_i = p + q_i(\mathrm{d}p/\mathrm{d}Q) - m = 0$. Rearranging terms, we obtain the optimality condition that

$$MR = p + q_i \frac{\mathrm{d}p}{\mathrm{d}Q} = m = MC, \tag{1.2}$$

where MR is the marginal revenue.

We obtain the Nash-Cournot equilibrium quantities by solving the n optimality equations for $q_1, q_2, \ldots, q_n$. Because the firms face the same marginal cost, they produce identical quantities, $q_1 = \cdots = q_n \equiv q$, in equilibrium so that industry output is $Q \equiv nq$. By substituting these identities into Equation (1.2), we can rewrite this expression as

$$L \equiv \frac{p - m}{p} = -\frac{1}{n\varepsilon} = -\frac{s}{\varepsilon}, \tag{1.3}$$

where $\varepsilon = (\mathrm{d}Q/\mathrm{d}p)(p/Q)$ is the market demand elasticity and $s \equiv q/Q = 1/n$ is the share of each firm. According to Equation (1.3), in the degenerate case in which $n = 1$, we find the usual result that a monopoly price–cost markup equals the negative of one over the market elasticity. Where n is greater than or equal to two, we learn that the markup equals the negative of one over $n\varepsilon$ or s/ε, which is the elasticity of demand facing each firm in equilibrium. As the number of firms grows without bound, we obtain the competitive result that the markup is zero.

Now, suppose that we generalize the optimality condition to allow the marginal cost to vary across firms, m_i. By summing with respect to i, Cowling and Waterson (1976) observed that the weighted average price–cost margin for the industry is

$$L = \sum_i s_i \frac{p - m_i}{p} = -\sum_i \frac{s_i^2}{\varepsilon} = -\frac{HHI}{\varepsilon}, \tag{1.4}$$

where $s_i = q_i/Q$ and HHI is the Herfindahl-Hirschman Index, the sum of the square of the share of each firm in the market is a measure of industry concentration. Thus, at least for these Cournot models, we obtain a clear relationship between Lerner's index, a measure of performance, and the structure of the market as captured by the number of firms, the share of each firm, or the HHI.

If the industry is monopolistically competitive and firms enter until the marginal firm earns zero profits, then n depends on the average cost function of the firms as well as the elasticity of demand. That is, both fixed and variable costs matter. Actions by governments or others that prevent firms from entering the industry (e.g., licensing laws, taxi medallions) similarly affect a firm's market power. Actions by a firm to differentiate physically its product or to convince consumers that its product is different through advertising, raise the elasticity of demand the firm faces and hence its market power. Thus, these "explanations" for market power can

be built directly into the demand curve, the cost curve, or a market equilibrium equation in a structural model.

STATIC MODELS

All the subsequent studies that we discuss throughout the book make use of these basic ideas to model the optimizing behavior of firms conditional on the demand, cost, and other explanations for market power. The modern static approach, which has been used for a third of a century, attempts to estimate the unobserved degree of market power and to determine its causes by estimating optimality and other equations simultaneously.[3]

There are three major types of modern static models, which are based on a single-period oligopoly model. In Chapter 3, we examine the first two types of studies, which examine industries that (are assumed to) produce homogeneous products. One set of studies uses structural models whereas the other uses comparative statics to estimate or test for market power with reduced-form or nonparametric models. In Chapter 4, we show how more recent studies extend the structural models to examine differentiated goods markets.

In Chapter 3, we discuss studies based on industry-level data. The first step is to construct each firm's first-order, or optimality, conditions via a profit maximization procedure similar to Equation (1.2). However, these models do not assume that the industry necessarily has a Cournot equilibrium. Rather, they use a more flexible approach that permits a large family of possible equilibria including monopoly, Cournot oligopoly, other oligopoly, and competitive outcomes. Typically, researchers use a single parameter to index the various possible outcomes. For example, they may introduce a "conjectural variation" parameter v that characterizes how Firm i believes its rivals will respond to a change in its output: $v = \mathrm{d}Q_{-i}/\mathrm{d}q_i$, where Q_{-i} is the collective output of all the firms except for Firm i. Depending on the value of v, firms could behave competitively, oligopolistically, or collusively, and hence v is a measure of market power.

Because the researchers do not know marginal cost, they usually express marginal cost as a function of output and possibly other variables such as factor prices. Researchers then aggregate these optimality conditions to get a market-level condition. Next, they simultaneously estimate that equation and other equations describing the market, such as the market demand equation. Their purpose is to simultaneously identify the market-power parameter and the marginal cost.

[3] Bresnahan (1989) summarized most of the modern static approaches through the 1980s. The June 1987 issue of the *Journal of Industrial Economics*, entitled "The Empirical Renaissance in Industrial Economics," edited by Bresnahan and Schmalensee, contains a number of the most important papers as of that time and a brief summary of the literature by the editors (Bresnahan and Schmalensee, 1987). See also Geroski, Phlips, and Ulph (1985), which surveys the literature on measuring conjectural variations and monopoly power.

This approach makes heavy use of assumptions about the exact specification of structural equations and may introduce constructs like conjectural variations. Consequently, other researchers prefer to use an approach than makes fewer assumptions. They employ reduced-form or nonparametric models to measure market power, bound it, or test whether the data are consistent with one market structure versus another. Because they relax the assumptions in the structural model, these researchers may not be able to estimate the degree of market power, but they can test the null hypothesis that the market is competitive. Typically, underlying the reduced-form approach is a comparative statics analysis. One might choose this simpler approach rather than the richer structural models because of limits on data or as a way to avoid the danger of specification bias caused by incorrect choice of functional form.

In Chapter 4, we turn to more recent static models in which researchers focus on individual firms rather than markets and relax the assumption that all firms produce identical products. The emphasis in most of these generalized models is on the role of factors that affect demand in determining market power. Most studies use one of two approaches to estimating market power with multiple firms with differentiated products. The first approach is based on estimating the firms' residual demand functions. In the second approach, researchers simultaneous estimate the complete demand system and the set of optimality conditions. Both approaches are based on a structural model and generalize the homogeneous good, one-sector model. Several of the earliest studies derive residual demand functions for individual firms from the full structural model. For example, the inverse residual demand curve facing Firm i, $d_i(q_i)$, is the inverse market demand curve, $p(Q)$, less the supply curve, $S_o(p)$, of all other firms (where it is defined):

$$d_i(q_i) = p(Q) - S_o(p). \tag{1.5}$$

Regardless of market structure, we again find that Firm i maximizes its profit by equating its marginal revenue (corresponding to its residual demand curve) to its marginal cost,

$$MR_i = p_i\left(1 + \frac{1}{\varepsilon_i}\right) = MC_i, \tag{1.6}$$

where ε_i is the elasticity of its residual demand curve, Equation (1.5). By rearranging this expression in the usual way, we obtain the Lerner index expression:

$$L_i = \frac{p_i - MC_i}{p_i} = -\frac{1}{\varepsilon_i}. \tag{1.7}$$

The markup depends solely on the elasticity of residual demand. From Equation (1.7), it appears that we can calculate the degree of market power, L_i, without using cost information; however, we need to use information about at least other firms' costs to determine the residual demand curve properly. Thus, these studies

take into account both demand and cost factors, and estimate a market-power or conjectural variation parameter directly.

More recent studies estimate full systems of demand and optimality equations. Most of these studies use one of two approaches to specify and estimate the demand system. In the neoclassical demand-system approach, a demand model is estimated based on market-level data using a flexible functional form and imposing restrictions from economic theory such as adding up, symmetry, and homogeneity properties.

In the most recent work, researchers use market share or consumer-level data to estimate the consumer demand system, typically without imposing restrictions from economic theory. The best-known studies using this approach employ a random utility model. Combining these more sophisticated demand estimates with optimality conditions, researchers determine market power using similar methods to those used in the earlier static structural models.

DYNAMICS

In Chapters 5 through 8, we study the generalization of the early static structural models to dynamic models. Although in these dynamic models the emphasis is on estimating market power, they can also provide an answer to the second question on the determinants of market power.

In these models, firms interact over many periods – they play a dynamic game. Each firm maximizes its expected present discounted value of the stream of its future profits. We want to use observations on firms' behavior to measure the extent of competition in a dynamic game. The estimation problem depends on the reason for the dynamics, the type of game that firms play, and on the data observed by the researcher. We distinguish between *strategic* and *fundamental* reasons why firms might play a dynamic game rather than engage in a sequence of static games.

It is possible that firms interact in a market repeatedly over time, where a firm's decision this period does not affect demand, costs, or other variables that affect profit in subsequent periods. Nevertheless, a firm's current decision might alter its rivals' beliefs about how that firm will behave in the future, thus altering the rivals' equilibrium future behavior. If dynamic interactions arise because firms think that their rivals will respond in the future to their current action, because of a change in beliefs, we say that the reason for the dynamics is strategic.

In contrast, a firm's decision about advertising, capital investment, or other long-lasting factors in the current period might directly affect future demand or cost. In these cases, current decisions alter a payoff-relevant state variable such as the stock of loyal customers, or the stock of capital. The change in this market "fundamental" alters future decisions by all firms. In this case, we describe the source of dynamics as fundamental.

With both strategic and fundamental dynamics, current actions affect future actions, but the mechanism that creates this dynamic relation is different, and it

calls for different estimation methods. Although dynamics may be important for both strategic and fundamental reasons, we discuss them separately. Chapter 5 considers the estimation problem when the only source of dynamics is strategic. Chapters 6 through 8 discuss models used when there is a fundamental source of dynamics.

Chapter 5 reviews basic concepts from game theory, including subgame perfection and the Folk theorem of supergames, and explains their role in the estimation problem. A supergame is a static one-shot game that (with positive probability) is repeated infinitely often. The simplest way to think of a subgame is to envision a game tree with infinitely many nodes. In each period, players begin at a particular node. Their collective decisions in that period send them to a new node in the next period. The game beginning at each node is referred to as a "subgame." A strategy is a mapping from a player's information set to the player's set of actions; in this sense, a strategy "instructs" each agent how to behave in every subgame. An equilibrium strategy is a collection of strategies; one for each player, for which each player's strategy is a best response to the rivals' strategies. An equilibrium is "subgame perfect" if and only if the set of strategies is an equilibrium for each possible subgame (i.e., at each node) – not just at the subgames that are actually visited along the equilibrium trajectory. The Folk theorem shows that there are (typically) infinitely many subgame perfect equilibria. This result is empirically important because it means that we (typically) cannot rely on economic theory to identify a unique equilibrium in dynamic settings.

Chapter 5 considers two classes of models, using either punishment strategies or continuous strategies. The equilibrium strategy involving punishments typically requires that players use one action (such as choosing a price or a quantity) in the cooperative phase, and a different action in the punishment phase. The outcome in the cooperative phase leads to a higher payoff than in the noncooperative equilibrium to the one-shot game, but it need not produce the cartel solution.

The use of the term "cooperative" is potentially confusing because the underlying game is noncooperative. The noncooperative equilibrium consists of a set of strategies, where each firm's strategy is a best response to other firms' equilibrium strategies. The (credible) threat of punishment causes firms to resist the temptation to deviate from their cooperative action. Actions in the cooperative phase might look approximately or even exactly like cooperation, but these actions are sustained as part of a noncooperative equilibrium. In this setting, even if the cooperative phase produces the cartel solution, firms have not colluded. In other words, noncooperative behavior in a supergame is consistent with a wide range of outcomes, sometimes extending from the one-shot noncooperative equilibrium to the cartel solution. The econometric objective is to determine where the outcome lies on this continuum of possible equilibria.

Models with continuous strategies assume that agents use a decision rule that depends continuously on their rivals' previous actions. In contrast, with the punishment strategies described earlier, the observed actions (in most models) take

one of two values, the cooperative and the noncooperative levels. The two types of models are similar, however, in that both assume that current actions affect rivals' future actions, even though current actions do not alter an economic fundamental (such as demand or costs).

Chapter 6 provides an overview of the concepts needed to estimate dynamic models based on fundamental rather than purely strategic considerations. We illustrate the role of fundamentals using both production-driven and demand-driven scenarios. For example, current investment in physical capital affects future production costs, and current advertising affects the level of future demand. In these kinds of models, firms make decisions that may affect their current profit only (a "static decision") in addition to decisions that affect their future stream of profit (a "dynamic decision"). The level of output in the current period (the static decision) affects current profits, but need have no direct effect on future profit; in contrast, the current level of investment (the dynamic decision) does have a direct effect on future profits. Empirical estimates of market power attempt to measure how close an industry is to cartel or cooperative behavior. However, firms might have different degrees of cooperation regarding their static and dynamic decisions. If the game is misspecified, estimates of a measure of the degree of cooperation with respect to one type of decision may be biased. In addition, increased cooperation with respect to one type of decision can have ambiguous effects on firms' profits and on social welfare. Thus, the coexistence of static and dynamic decisions complicates both the problem of obtaining reasonable estimates of market power, and also interpreting the welfare implications of those measures.

There are many types of dynamic strategies. We explain the difference between two important types of strategies, *open-loop* and *Markov Perfect*, using a deterministic setting in which firms incur convex adjustment costs to change the level of their capital stock. With Markov strategies, firms understand that their current actions affect a state variable (e.g., a level of capital or a stock of loyal customers) that affects rivals' future actions. Firms take their rivals' strategies (i.e., their decision rules) as given, and understand that by altering the state variable they can affect rivals' future actions. In contrast, with open-loop strategies, firms believe that their rivals' strategies do not depend on these state variables. Firms therefore behave as if their actions have no effect on their rivals' future actions. The open-loop equilibrium can be obtained by solving a one-agent optimal control problem, but finding the Markov Perfect equilibrium requires the solution to a game.

When agents have symmetric power (as distinct from a leader–follower role), both the open-loop and the Markov Perfect equilibria are "time-consistent." To understand time consistency, suppose that we compute the equilibrium at an arbitrary time – for example, time 0 – and assume that firms follow their equilibrium strategies until a future time, t. Because (by assumption) firms have followed the equilibrium strategies up to time t, the game is "on the equilibrium trajectory" at time t. Time consistency of the equilibrium means that the equilibrium computed at time t equals the continuation of the equilibrium that was computed at time 0.

Subgame perfection is a stronger condition; it requires that the continuation of the initial equilibrium is an equilibrium starting at *any* subgame, not just the subgames that are visited in equilibrium. That is, the equilibrium decision rules computed at time 0 are also equilibrium decision rules at time t even if one (or more) firm deviated from its equilibrium rule between the two points in time (thus causing the state variable to be "off the equilibrium path"). The Markov Perfect equilibrium is a particular member of the set of subgame perfect equilibria. The open-loop equilibrium is, in general, not subgame perfect. For this reason, the Markov Perfect equilibrium is usually considered to be a more plausible solution concept than the open-loop equilibrium.

There typically are many Markov Perfect equilibria. In some cases, a "good" outcome can be supported by credible threats to adopt punitive strategies following "bad" behavior. This kind of strategy is reminiscent of the punishment strategies that are used in supergames; however, in supergame models, the current actions are explicit functions of past actions, whereas in a Markov equilibrium, the current action depends only on a directly payoff-relevant variable, such as current capital stocks. Non-uniqueness of Markov equilibria also arises for a more subtle reason; the necessary (equilibrium) conditions do not pin down agents' beliefs about how rivals will respond to deviations from equilibrium in a steady state. This reason for non-uniqueness is similar to the reason why there are multiple consistent conjectural variation equilibria in static games.

Chapter 7 discusses the estimation of dynamic models under different assumptions about the equilibrium: open-loop, Markov Perfect, or a hybrid. The chapter describes in detail two models that differ in the fundamental reasons for dynamics. These models are used to address different empirical questions. The first of these, the sticky price model, provides a straightforward dynamic generalization of the static conjectural variation models used in the first part of this book. We discuss the estimation of this model under the assumption that firms use open-loop strategies. A conjectural variations parameter provides an index that measures the extent to which the observed outcome is closer to a competitive or collusive equilibrium. The necessary condition for optimality, the Euler Equation, is the basic estimation equation, just as the first-order conditions to static optimization problems are the basis for estimating static models.

The second model in Chapter 7 studies firms' advertising decisions. Here we assume that the equilibrium is Nash-in-advertising. We compare the necessary conditions (the Euler equations) when firms use either an open-loop rule or a Markov Perfect rule, and we consider a hybrid that lies "between" these two types of equilibria. By estimating this hybrid model, it is possible to determine whether firms use a naïve (open-loop) or a more sophisticated strategy in which they recognize that rivals may respond to changes in the endogenous state variable.

The final section of this chapter discusses a relatively new technique for estimating the parameters of the objective function in a game in which firms use Markov Perfect strategies. Using an example, we show how this estimation strategy can be