

Oligopoly Dynamics

Luís Cabral

New York University, IESE and CEPR

This draft: December 2011

Abstract. I argue that dynamic oligopoly models are an area of industrial organization where much work needs to be done and much work can be done. In some particular settings (e.g., network industries), dynamic oligopoly models provide sensible answers when static models fall short of doing so.

Presidential Address delivered at the 38th Annual Conference of the European Association for Research in Industrial Economics (Stockholm, September 2011). Luís Cabral is W R Berkley Term Professor of Economics, Stern School of Business, New York University; Academic Director, IESE Business School's New York Center; SP-SP Research Fellow, IESE; and Research Fellow, CEPR (London); lcabral@stern.nyu.edu. He served as EARIE president from 2009–2011.

1. Introduction

The field of industrial organization has undergone several “revolutions.” The 1980s, for example, witnessed the “game theory revolution,” the systematic application of (relatively novel) game theoretic concepts to many problems of strategic behavior. A few years later, a series of authors attempted to model strategic behavior empirically, replacing old reduced form regressions with carefully crafted structural econometric models: the so-called “New Industrial Organization” movement was under way.

What will be the next “revolution” in industrial organization: Behavioral economics? Numerical methods? Field and laboratory experiments? While it is hard if not impossible to make such prediction, in this paper I will argue that dynamic oligopoly models are an area where much work needs to be done and much work can be done. In particular, I will argue that, in some settings, dynamic oligopoly models provide considerable value added with respect to static models.

I will proceed as follows. In the next section I propose a simple classification of dynamic models. Then, I lay down a particular framework for the study of dynamic oligopoly, in particular oligopoly industries with network effects. The core of my argument is developed in Section 4, where I present two applications where a dynamic model provides a sensible answer whereas a static falls short of doing so. Section 5 concludes the paper.

2. A taxonomy of dynamic oligopoly models

There are many types of dynamic oligopoly models. For the purposes of this paper, I have a particular type in mind. I will thus attempt a brief classification and indicate what I mean by dynamic oligopoly models in the present context.

In many situations, time is modeled by an extensive form game with multiple stages. For example, Klemperer (1987) models the dynamics of markets with switching costs by means of a two-period model: in the first period, consumers are not locked-in to any seller, whereas in the second period they are locked-in to the seller they purchased from in the first period. While multi-period models of this type have dynamics in them, they suffer from a potentially serious limitation, namely the bias introduced by there being a first and a last period (in real-world markets, such beginnings and endings are hard to determine).

A solution to the end-of-the-world problem of two-period models is to consider an infinitely repeated game.¹ For example, Friedman (1971) offers the first model of collusion in dynamic oligopoly competition: he shows that grim strategies with reversion from monopoly output to Cournot output form an equilibrium if the discount factor is sufficiently high. Models like these are dynamic in the sense that strategies are history dependent, and thus time plays an important role. However, the stage game that induces the repeated game is essentially a static game. In other words, other than history-dependent strategies, there is no “physical” connection between periods.

This finally leads me to what I believe is a properly defined dynamic oligopoly model: a model where time extends indefinitely (as in a repeated game) but where there are state variables providing a “physical” link across periods. The past two decades or so have witnessed a considerable increase in the interest for this type of models. For all their

1. Infinitely repeated games are commonly understood as indefinitely repeated games, but somehow the term “infinitely” stuck.

variety, I believe there is an interesting contrast between what I call “investment” models and what I call “pricing” models.

Figure 1 summarizes the main features of these two types of models. Investment models, which correspond to the left panel, are characterized by a state variable ω_{it} such as product quality, production capacity, or cost. The critical dynamic element is given by an investment function $x_{it} \rightarrow \Delta \omega_{it}$ where x_{it} is firm i 's investment at time t and $\Delta \omega_{it}$ the resulting change in firm i 's state. Typically, in these models demand and product market competition are treated in a non-dynamic manner, that is, each period product market competition takes place given the state of the game. One of the earliest examples of this type of models is given by Ericson and Pakes (1995); a good number of models have been developed along these lines (see Doraszelski and Pakes, 2007, for an excellent survey).

In pricing models, by contrast, the critical dynamic element corresponds to the firms' pricing decisions and the impact it has on market shares. In other words, the state space is defined not so much by the firms' “physical” attributes as it is by their market shares: the number of consumers who are attached to the firm. Beggs and Klemperer (1990) provides one of the earliest examples of this type of model. Because of switching costs, consumers are “attached” to a given firm (the firm they bought from last time). In each period a measure of consumers enters the market anew and firms compete for their custom.

Many models, especially more recent ones, include features of investment and pricing models alike. Even then, I believe the classification as investment or pricing models, stylized as it may be, can prove useful. For example, in Cabral and Riordan's (1994) model of dynamic pricing with learning curves the state space is defined by each firm's cumulative sales, which determine their cost. To the extent that the state space includes firm's cost, the model might be thought of as an investment model. However, the critical dynamic element is given by the pricing decision and how it affects current demand, and consequently the change in state space. In this sense, the model is essentially a pricing model.

In this paper, I will focus primarily on this type of models, that is, models where prices affect current demand and current demand changes the state space.

3. A framework for dynamic analysis

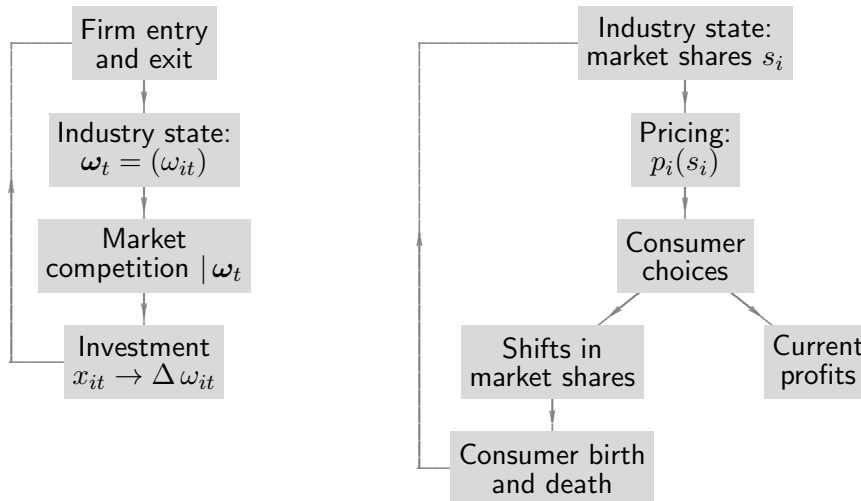
Consider an industry which exists in continuous time and lasts into the indefinite future. Both firms and consumers are long lived. Consumers must make some durable decision infrequently and some non-durable decisions constantly. Examples of durable decisions include the choice of (a) a telephone network, (b) an operating system, (c) a credit card; corresponding examples of non-durable decisions are (a) the number of calls to make, (b) applications software to purchase and upgrade, (c) purchases to make with a credit card.²

Specifically, I assume that, for each consumer, a moment arrives when he needs to reassess his durable good choice. I don't explicitly model the reason for such reassessment event (it could be, for example, that the consumer's phone was lost or that he changed jobs); I simply model it as a Poisson process with arrival rate λ . I moreover assume that, in-between these moments, the consumer sticks by his previous durable decision (using a certain phone or a certain operating system) without questioning it. Effectively, it's as if the

2. Some of the decisions I call “non-durable,” such as software purchases, may well be durable. The important feature is that “durable” decisions are more durable than “non-durable” decisions.

Figure 1

Dynamic models: investment models (left) and pricing models (right).



consumer's switching cost alternated between infinite (most of the time) and zero (selected moments which arrive as a Poisson process).

Next I discretize time such that, in each period, there is exactly one consumer who reassesses his durable decision. Essentially, I consider the expected time between two consecutive industry reassessment moments (which typically will correspond to two different consumers) as a period in my discrete time model. By assuming risk-neutral agents, I derive the discount factor of this discretized models as $\delta = \exp(-r/\lambda)$, where r is the continuous time discount rate. Note that, in the discrete-time model, the *actual* duration of each period varies, just as the arrival of a Poisson event varies. However, for the purpose of computing value functions the fact that agents are risk neutral allows me to work with expected period duration. In other words, the fixed duration of a period in the discrete-time model should not be interpreted literally, rather as a reduced form (average duration) of a real-time Poisson process. Similarly, the assumption that there is exactly one reassessment per period, which at first may seem rather odd, is in fact the natural implication of the discretization strategy I follow: by construction there is exactly one event between two consecutive Poisson events.

In each period, the state of the game is given by the number of consumers "assigned" to each firm, that is, the number of consumers who have made a durable decision to buy from each firm. I consider Markov equilibria, whereby firm and consumer decisions are a function of the the state of the game.

Firm strategies consist of pricing. There are many possible assumptions regarding the prices set by firms. For the purpose of this paper, I consider the case when firms set a price for the durable good (which is therefore only paid by the consumers who make a durable decision), as well as a price for the nondurable good (which is paid by all of the consumers assigned to a given firm).

Consumers, in turn, must decide how much to purchase the nondurable good in each period; moreover, some consumers (one per period) are asked to make the durable decision of which firm they want to be attached to. There are several possible assumptions regarding

consumer heterogeneity (beyond the fact that they are attached to a particular firm). For the purpose of this paper, I assume that consumers draw firm-specific “stand-alone” utility values (as well as the value of the outside option) each time they have to make a durable decision. Moreover, I assume those firm-specific utility values are normally distributed and that the outside option has value $-\infty$, that is, a consumer always chooses to be with one network. Finally, the utility consumer derive from using their network is an increasing function of the number of other consumers using the same network (network effects). In sum, I assume consumer heterogeneity in the “stand-alone” value of a network and homogeneity in the network portion of consumer benefit.

The Markov equilibrium of this game generates a stochastic process of prices and market shares, which in turn leads to a series of stationary distributions, in particular a stationary distribution of market shares. The sources of randomness are Nature’s choice of the next consumer to make a durable decision as well as the utility values generated at that point.³

4. Applications

Deriving and solving a dynamic model of the sort considered in the previous section is considerably more burdensome than the corresponding static model. Not only do we need to solve for the firm value functions, we must also solve for the consumer value functions. Is it worth going through all of this effort? In this section, I consider two questions where a dynamic model is clearly required. In some cases static models produce different answers than dynamic models. In the cases considered in this section, a static model would produce no answer — or no meaningful answer, I would claim — to the questions at hand.

■ **Network effects as a barrier to entry.** There is a considerable debate, dating back to Bain and Stigler, over the definition of barriers to entry.⁴ I propose a specific measure of barriers to entry, the evaluation of which requires a dynamic model.⁵ I suggest that network effects create a barrier to entry to the extent that a new entrant’s value is lower than it would be if there were no network effects. Suppose that the parameter ψ measures the intensity of network effects. Specifically, the utility that a consumer derives in each period is given by a firm-specific utility value plus ψn , where n is the size of the network the consumer belongs to. One way to measure the decrease in entrant’s value due to network effects is to compute

$$B = \frac{v(0) |_{\psi=0} - v(0) |_{\psi>0}}{v(0) |_{\psi=0}} \quad (1)$$

where $v(0) |_{\psi>0}$ is the the value function of an entrant (that is, a firm with a network of size zero) when network effects exist ($\psi > 0$); whereas $v(0) |_{\psi=0}$ is the corresponding value function when network effects are absent.⁶ Notice that neither the Bain nor the

3. A more detailed description of this framework may be found in Cabral (2011).


4. See Cabral (2008) for a discussion on the concept of barriers to entry. See also the 2004 AEA session on barriers to entry, including the papers by Carlton (2004); McAfee, Mialon and Williams (2004); and Schmalensee (2004).

5. See Jullien (2001), Llobet and Manove (2006) for different perspectives on the implications of network effects for entry.

6. An alternative measure would be difference between the incumbent’s value, $v(100)$, and an entrant’s value, $v(0)$. This alternative is closer to Gilbert’s (1989) notion of barriers to entry (“a rent that is derived from incumbency”).

Figure 2

Network effects as a barrier to entry.



Stigler definitions would indicate that network effects create a barrier to entry.⁷ This is not entirely surprising: the Bain and Stigler definitions are essentially static, whereas the barrier to entry created by network effects is essentially a dynamic problem (and thus requires a dynamic model for its estimation).

It can be shown (Cabral, 2011) that $v(0) |_{\psi=0} > v(0) |_{\psi>0}$. In other words, the above measure B has positive sign, that is, network effects imply a barrier to entry. In order to go beyond that, that is, in order to estimate the magnitude of the barrier to entry, it is necessary to estimate or calibrate the model so as to obtain the relevant parameter values.

Accordingly, I construct a measure of network “die-hard” fans. Consider the following experiment: a new consumer must choose between a network of size zero and a network of size η . Suppose both networks set the same entry price. Suppose moreover that the new consumer is myopic, that is, assumes the current network size will persist into the indefinite future. A die-hard fan is one who chooses a network of size zero in these circumstances.

Suppose there are $n = 100$ potential adopters and that the difference in firm-specific utility values has a standardized normal distribution. It can be shown that, as $\psi = 0, .1, .2, .3$, the measure of die-hard fans is given by 50, 17.9, 3.3, and .3%, respectively. This gives an idea of what “reasonable” values of ψ might be.

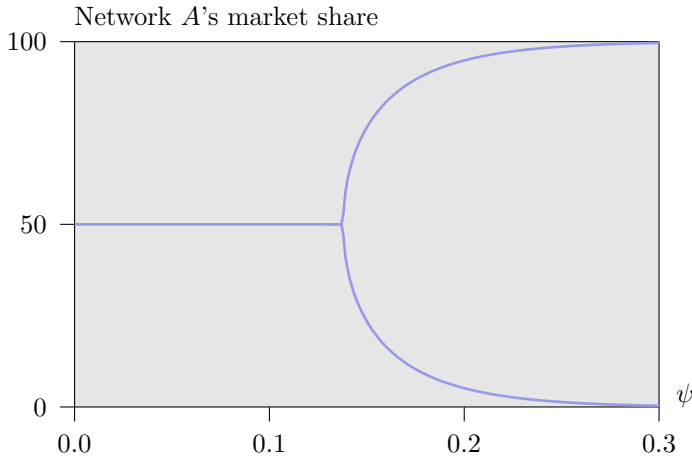
By further assuming a discount factor $\delta = 0.9$ and solving the model, I am able to compute $v(0)$ when $\psi = 0$ and $v(0)$ when $\psi > 0$. By comparing these values, I can compute the barrier to entry created by network effects. Specifically, Figure 2 plots the value of B , the entry barrier created by network effects, as a function of the degree of network effects, ψ . The figure suggests that even for “reasonable” values of the degree of network effects the barrier to entry created by network effects is substantial. For example, if the measure of die-hard fans is about 20%, then network effects imply a 50% cut in the value of an entrant.

How essential is the dynamic model for estimating this barrier to entry? In principle, one could have a static model where the incumbent has some installed base, a fraction of

7. Bain (1956) defined an entry barrier as the set of technology or product conditions that allow incumbent firms to earn economic profits in the long run. Stigler (1968) in turn offered an alternative definition: a cost of producing which must be borne by an entrant but not by an incumbent. See Gilbert (1989), Cabral (2008).

Figure 3

Equilibrium average market share in static model when both networks are equal quality.



which is locked in, and then derive the value of B exactly as in (1). In other words, the above definition of barriers to entry can be applied to static models as well: all we have to do is compare the entrant's equilibrium payoff with and without the putative source of entry barrier. The static model in question would then be thought of as a reduced form of a dynamic model where, following entry, some consumers might switch to the entrant while others don't. But there is a considerable degree of arbitrariness in the modeling of consumer behavior and installed base. Specifically, the dynamic model makes the (reasonable) assumption that, ultimately, all consumers currently attached to the incumbent are switchable: it's only a matter of time and pricing before they switch to the entrant. Qualitatively, one can model this as a distribution of switching costs in a static model, which in turn results in a positive barrier to entry. So, while qualitatively the static model mimics the dynamic model in showing that network effects create a barrier to entry, when it comes to estimating the actual value of such barrier the static model is of little help.

■ **Product improvement.** In the preceding analysis, I assumed that the two networks are of equal quality (though possibly of different sizes). In other words, each time a consumer reassesses his durable choice Nature generates each network's "stand-alone" utilities from the same distribution. To this each consumer adds utility by using the network, and this utility is a function of network size.

Suppose now that one of the networks is of superior quality, in the sense that network A 's "stand-alone" value is derived from a distribution with higher mean than network B 's. What does this imply for equilibrium network size? If network effects are absent, then the answer is quite straightforward. Dynamics don't really matter that much, and so a static model does the job. Specifically, consider the following static model: η consumers simultaneously receive preference for each of the networks and simultaneously decide which network to choose. In this case, if the two networks are of equal quality, then there is a unique equilibrium whereby each network's market share is on average 50%. If one of the networks improves its product then it's expected market share increases accordingly.

When network effects are present, however, then the question becomes considerably more complex. Figure 3 plots the equilibrium values of network A 's market share for each value

of the network effects parameter, ψ , assuming that the two network are of equal quality. As mentioned in the previous paragraph, if there are no network effects then there is a unique equilibrium with network A 's expected market share set at 50%. If however ψ is sufficiently high, then the equilibrium set bifurcates and two asymmetric equilibria emerge (there is also a symmetric, unstable, equilibrium). This creates a problem of equilibrium selection, which in turn creates a problem for comparative statics: how can one predict the impact of an improvement in network A 's product if one doesn't even know which equilibrium is being played? And even if we know which equilibrium is initially being played — for example, with network A taking the low market share — how do we know whether the increase in network A 's quality will tip the outcome towards the other equilibrium?

By contrast to the static model, the dynamic model provides a sensible answer to these questions. Figure 4 describes the case of product improvement with network effects set at $\psi = .2$. Consider first the case when both networks are of equal quality. Then the stationary distribution of network A 's market share is given by the blue line. The distribution is symmetric — as expected, given that the two networks are of equal quality. However, the distribution is bi-modal, meaning that, in the long-run, the system will stay most of the time in an asymmetric state, with one network of bigger size than the other one. This is roughly consistent with the prediction from the static model: as Figure 3 suggests, when $\psi = .2$ there are two asymmetric equilibria, where one of the networks is greater than the other one.

But the differences between static and dynamic model are greater than the similarities. First, I note that the degree of asymmetry predicted by the static model is considerably higher than the dynamic model: for $\psi = 2$, for example, the asymmetric equilibria shown in Figure 3 are farther apart than the modes of the stationary distribution in Figure 4. More important, whereas we have two asymmetric equilibria in the static model, the dynamic model implies a unique symmetric equilibrium, though this unique symmetric equilibrium implies asymmetric outcomes with very high probability.

The distinction between multiple, asymmetric equilibria and a unique symmetric equilibrium with multiple, asymmetric outcomes is not purely semantic. In particular, whereas the static model is silent when it comes to comparative statics, the dynamic model implies clearly defined comparative statics. As Figure 4 illustrates, an increase network A 's quality shifts the stationary distribution of market shares in network A 's favor. Moreover, such shift takes place in a continuous way. This allows for a very precise answer to the questions of the type, How does network A 's value change as a result of a quality improvement: all we need to do is to integrate over the distribution of possible outcomes implied by the unique equilibrium.

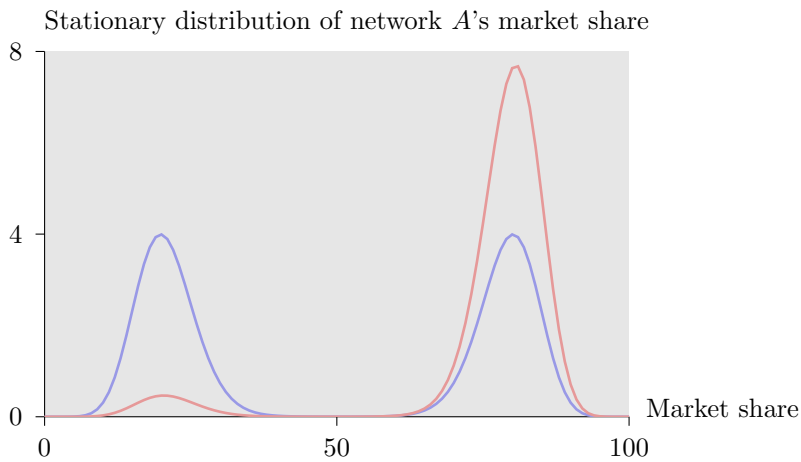
5. Discussion and final remarks

For many — maybe even for most — applications, static models provide a sufficient approximation of real-work markets and a workable comparative statics apparatus. In other instances, however, static models miss much of the action, perhaps even the central part of the action. This is when dynamic models can be so helpful.

Dynamic models are typically much more difficult to solve (which may partly explain why they are not more frequently used): even relatively simple dynamic models have no closed-form analytical solution. In these cases, the best solution, I think, is a combination

Figure 4

Effect of network A 's product improvement (from blue to red).



of analytical and numerical methods. Specifically, while no analytical solution exists, Taylor expansions induce linear systems that can be solved analytically. Although this approach restricts analysis to a measure-zero set of parameter values (around the point of expansion), it is a useful way of obtaining results and intuitions. Numerical methods can then be used to extend the results to a wider range of parameter values. Numerical methods have an added advantage: the repetition of numerical patterns suggests the possibility of general analytical results. We thus have an iterative and interactive process where the whole (analytical and numerical methods) is clearly greater than the sum of the parts.

References

- BEGGS, ALAN W, AND PAUL KLEMPERER (1992), “Multi-period Competition with Switching Costs,” *Econometrica* **60**, 651–666.
- CABRAL, LUÍS (2008), “Barriers to Entry,” in Steven N. Durlauf and Lawrence E. Blume (Eds), *The New Palgrave Dictionary of Economics, 2nd edition*, Basingstoke and New York: Palgrave Macmillan.
- CABRAL, LUÍS (2011), “Dynamic Price Competition with Network Effects,” *Review of Economic Studies* **78**, 83–111.
- CABRAL, LUÍS, AND MICHAEL H RIORDAN (1994), “The Learning Curve, Market Dominance and Predatory Pricing,” *Econometrica* **62**, 1115–1140.
- CARLTON, DENNIS W. (2004), “Why Barriers to Entry are Barriers to Understanding,” *American Economic Review (Papers and Proceedings)* **94** (2), 466–470.
- DORASZELSKI, ULRICH, AND ARIEL PAKES (2007), “A Framework for Applied Dynamic Analysis in IO,” in Mark Armstrong and Robert Porter (eds.), *Handbook of Industrial Organization*, Volume 3, North-Holland, Amsterdam.
- ERICSON, RICHARD AND ARIEL PAKES (1989), “Markov-Perfect Industry Dynamics: A Framework for Empirical Work,” *Review of Economics Studies* **62**, 53–82.
- JULLIEN, BRUNO (2001), “Competing in Network Industries: Divide and Conquer,” CEPR Discussion Paper No. 2883, July.
- LLOBET, GERARD, AND MICHAEL MANOVE (2006), “Network size and Network capture,” Working Paper, CEMFI and Boston University.
- MCAFEE, R. PRESTON; MIALON, HUGO M. AND WILLIAMS, MICHAEL A. (2004), “What Is a Barrier to Entry?,” *American Economic Review (Papers and Proceedings)* **94** (2), 461–465.
- SCHMALENSEE, RICHARD (2004), “Sunk Costs and Antitrust Barriers to Entry,” *American Economic Review (Papers and Proceedings)* **94** (2), 471–475.