2
Non-Competitive Markets

Nikolaos Georgantzís and Giuseppe Attanasi

Introduction

A market is not competitive when the agents acting in such a market have the power to influence the price, directly or indirectly, something that does not occur under perfect competition. Generally, these agents have market power because they are few in number, have access to relevant information and can foresee the interdependence between their strategies and those of others.

Among all the paradigms in economic theory, the theoretical predictions of oligopoly were the first to be examined in the laboratory. In the origins of experimental economics one can find the works of Chamberlin (1948) and of Smith (1962, 1964), who designed experiments to study a market with few agents that could reach the competitive equilibrium (see Chapters 1).

In this chapter, instead of surveying all experiments with few sellers,[1] we will adopt a narrower definition of the term "*oligopoly*," and will focus on experiments that were directly inspired by the basic oligopolistic models of Cournot (1838), Bertrand (1883), Hotelling (1929), von Stackelberg (1934), and similar. We will omit, therefore, other experiments, such as those of Chamberlin and Smith, which were designed with the aim of testing the predictive power of the competitive equilibrium model.

Most of the experiments we consider in this chapter have been run in the last three decades.[2] This literature can be considered as a new wave of experimental work, aiming at representing basic oligopolistic markets and testing their properties. This work represents a systematic

P. Branas-Garza et al. (eds.), *Experimental Economics*

attempt to study a similar, but not identical, question to that tackled by Chamberlin (1948) and Smith (1962, 1964). While the latter compared the results in the laboratory with predictions of the *competitive equilibrium*, the series of experiments we review here compare observed behavior with the corresponding *oligopolistic equilibria*.

The chapter is divided into independent sections, which refer to different parts of oligopolistic theory: including monopoly and a number of extensions of the basic models that have been chosen with the aim of providing a representative overview of experimental findings in this area.

Monopoly, price competition and product differentiation

The simplest and most frequently corroborated hypothesis regarding price-setting behavior in markets argues that, by relaxing the usual assumptions, we may converge to the equilibrium price of Bertrand. That is:

- If there are (alternative assumptions): few agents, with no experience and limited cognitive ability, with insufficient information of the conditions of the market,
- is the equilibrium reached through learning by trial and error?

That is, whether the experimental subjects reach the price predicted by theoretical models that assume (usual assumptions) perfect information and infinite cognitive capacity.

The simplest case is that of *monopoly*. Assume that you are participating in an experiment in which you are a monopolist and you face a demand function unknown to you. In fact, this function could be, for instance, $Q = 100 - 2p$, which assumes that sales in your firm decrease when the price increases. Given that this demand function is unknown, you can only try to guess the price that maximizes your profits by testing what happens as you set different prices. For simplicity, let us assume that your cost is $C = 0$; thus, your payoff from the experiment coincides with your profits.

With an *unknown demand function*, you cannot calculate the optimal price, as you would do during a microeconomics class. As anticipated above, one way to calculate which price maximizes your profits as fast as possible is to set any price and observe the sales and the profits. Then you would change your price to see the reaction in sales and profits. Subsequently, you would use the results of these first two periods to choose the price for the third period, and so on. By accumulating information, you could start noticing how to get closer to the profit-maximizing price. Classroom experiments with college students have shown that the optimal price $p^* = 25$ is reached after six to ten trial periods.

Therefore, for a monopoly with no information on the demand function, laboratory experiments confirm that learning by trial and error leads to convergence on the perfect information prediction. This happens with few repetitions and without requiring specific cognitive abilities.

However, the result is different when some complexity in the problem facing the monopolist is introduced. For instance, suppose that, as a subject participating in the experiment, you are asked to simultaneously set the *prices of two products*.[3] To better understand the additional difficulty, consider a firm that aims to find the optimal prices for its two substitute products/services, such as two different flights to one destination operated by the same airline. In this framework, experiments have demonstrated a significant deviation in the strategies a subject plays with respect to the one predicted by the theory. The deviation lasts even after several attempts, unless the subject receives information on the cross effects of the demand of a product on that of the other product.[4]

In the real world monopolists frequently face decision-making problems that are even more complex, like when they have to set prices to manage a dynamic system: for instance, a renewable resource. García-Gallego *et al.* (2008) study a market of this sort, and find that subjects systematically fail to learn, and are unable to converge towards the optimal level of resource preservation, even when playing for 50 periods.

In summary, we can argue that human subjects can learn how to set the optimal price in a monopoly without the (extremely high) level of information and rationality assumed by the theoretical models in the textbooks. However, this is true only as long as the complexity of the environment does not exceed a certain limit. Above this limit, a trial-and-error learning procedure is not enough to allow experimental subjects to arrive at the monopolistic equilibrium price.

Let us now analyze the problem of price-setting in the case of an *oligopoly*. Fouraker and Siegel (1963) ran the first experiments on this topic. Their focus was on the importance of the type of information transmitted between subjects playing in the role of firms in an oligopolistic market. Their objective was to evaluate the predictive capacity of the theoretical oligopolistic equilibrium. Their experiments confirm two results of enormous interest:

i. When subjects receive private information on their own profits, they tend to converge towards the Nash equilibrium prices (Bertrand-Nash equilibrium).
ii. If subjects also have available information on the profits of their competitors, then they will set prices higher than those predicted by the non-cooperative Nash equilibrium: their behavior shows a certain level of collusion.

The level of *collusion* dramatically increases if experimental subjects are given the possibility to communicate among themselves. Fonseca and Normann (2012) compare pricing behavior with and without the possibility of communicating between firms in Bertrand oligopolies with various numbers of firms. They find strong evidence that *communication* helps to obtain higher profits for any number of firms in the market: communication helps firms coordinate in collusive pricing schemes. However, the gain from communicating is non-monotonic in the number of firms, with medium-sized industries having the largest additional profit from communication. They also find that industries continue to collude successfully even after communication is prevented.

Let us now focus on the *demand side*: the most recent oligopoly experiments have adopted mechanisms of market clearing where the buyers' behavior is simulated with continuous demand functions.[5] For instance, García-Gallego (1998) uses a system of symmetric demand functions, composed by *n* equations as the following:

$$q_i = a - b \cdot p_i + \theta \cdot \sum_{j \neq i} p_j$$

where p_i is the price of the variety i, n is the number of varieties available in the market, (each variety offered by a different firm), while j represents each of the substitutes of variety i.

The function tells us that the quantity demanded q_i of the good of variety i decreases as its price increases (b is a positive number), but it also increases as the price of any other variety increases. The parameter θ indicates the level of interdependency between varieties, in this case between variety i and the other $n - 1$ varieties. Such a demand system, where θ is positive and lower than b, corresponds to the case where each variety can be imperfectly substituted by any other variety in the market. Further, the unitary cost c is assumed to be constant and the same for all varieties, and there are no fixed costs.

In García-Gallego's experiment (1998), even though subjects systematically attempt to tacitly coordinate in setting prices at a level above that of the non-cooperative equilibrium, it is found that the Bertrand-Nash equilibrium strongly attracts individual strategies, especially in the second half of the market periods. This work shows that a 35-period horizon is sufficient to allow most subjects to converge surprisingly closely on Bertrand's prediction, with some experimental sessions in which predicted and observed behavior coincide. Figure 2.1 shows the evolution of prices in a typical session of the experiment.

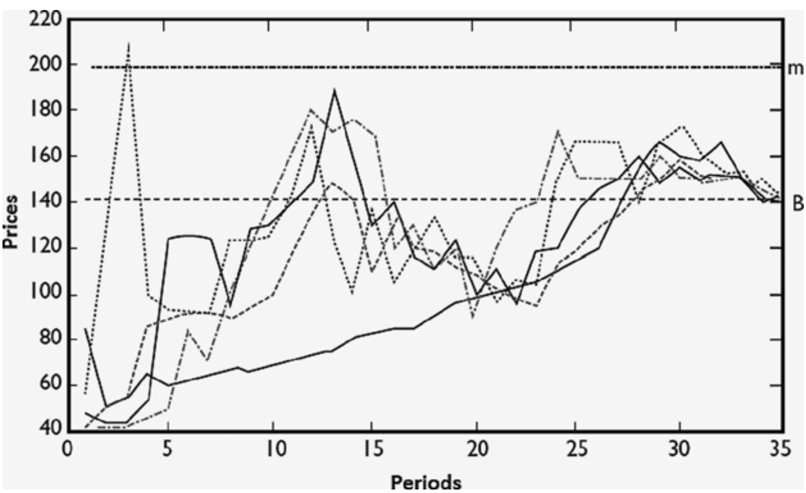

Figure 2.1 Bertrand-Nash vs. actual behavior in a differentiated oligopoly with five varieties

Note: After few periods of price volatility, subjects' strategies clearly converge to the Bertrand-Nash prediction (B) and move away from the monopoly (collusive) price (m).
Source: García-Gallego and Georgantzís (2001b)

Subsequently, García-Gallego and Georgantzís (2001a) implement the same conditions of demand and costs as those in García-Gallego (1998) to test the predictive power of Bertrand-Nash with *multi-product firms*. The theory predicts that, if the same multiproduct firm jointly produces two or more substitute goods, their prices will be higher than if independent competitors offered the goods. Although the problem of a multiproduct firm is much more complex than that of a firm with a single product, the existence of a multiproduct firm leads to the prediction of an asymmetric Bertrand equilibrium, where the firms producing more goods tend to set higher prices. Surprisingly, the experiments in García-Gallego and Georgantzís (2001a) suggest that multiproduct firms do not understand the strategic profits they can derive from their multiproduct market power. This is why they behave as if their products were competing with each other.

Therefore, trial-and-error learning in an oligopolistic market with uniproduct firms leads to strategies that are close to the Bertrand-Nash equilibrium. However, in the multiproduct case, trial-and-error learning does not support the corresponding asymmetric Bertrand-Nash equilibrium. Davis and Wilson (2005) have reinterpreted this result in terms of its consequences for mergers policy: if firms do not realize their market power, behavior after the *merger* can be as competitive as before the merger.[6]

However, García-Gallego and Georgantzís (2001a) also carry out a treatment imposing an exogenous norm that limits the multiproduct firm to change its own prices all in the same direction (*price parallelism*). In this treatment, subjects tend to adopt the limit strategies of the Bertrand-Nash multiproduct equilibrium (that is, close to the collusive equilibrium)[7]. Figure 2.2 presents examples of sessions with multiproduct oligopolies.

The reader can identify that the prediction of the Bertrand-Nash multiproduct equilibrium for a firm that jointly sets the prices of three varieties (B3 in Figure 2.2) is the highest, followed by the same prediction for two varieties (B2), and that both of them are greater than the Bertrand-Nash equilibrium with uniproduct firms (B). In Figure 2.2 on the left hand side, a typical session is shown where, thanks to the exogenous imposition of the price-parallelism norm, firms' prices converge to the collusive equilibrium. On the right hand side, a session is shown where, without such exogenous imposition,

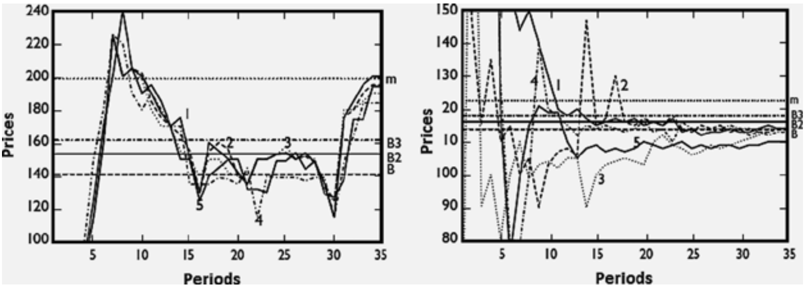

Figure 2.2 Multiproduct oligopoly

Note: Left: Convergence to the collusive equilibrium price (*m*) in a differentiated oligopoly with multiproduct firms, five varieties, and the price-parallelism norm. Right: Failure in reaching the collusive equilibrium price (*m*) because of no adoption of a price-parallelism norm.

Source: García-Gallego and Georgantzís (2001b)

multiproduct firms do not adopt price parallelism and cannot avoid converging to the Bertrand-Nash equilibrium with uniproduct firms.

Likewise, firms can affect the level of market competition by strategically choosing to differentiate one from the other. In the standard *product differentiation* model of Hotelling (1929), firms decide first on their *location*, which represents a variety in a continuous and closed product space, and then they compete in price-setting.

As in the case of several other phenomena for which it is very difficult to empirically test economic theories with real data, models with product differentiation have also been examined in the laboratory. Brown-Kruse and Schenk (2000), Collins and Sherstyuk (2000), and Huck *et al.* (2002b) have experimentally studied spatial markets with two, three and four firms, respectively.

These three works describe experiments in which the participants only choose the location of their firm, while the *prices are exogenously imposed* on them. In this context, two clear-cut theoretical predictions can be provided:

- There is a "minimal" product differentiation as in the non-cooperative equilibrium when subjects cannot communicate among them.
- There is an "intermediate" differentiation due to collusion when subjects can communicate between them.

Both theoretical predictions have been corroborated by the above-mentioned experimental studies. That is, firms that use location as the only strategic variable, in absence of communication tend to agglomerate in the middle of the product space (segment). Likewise, as we will see in Chapter 9, political parties tend to adopt ideological positions close to the median voter's preference. In the presence of communication, firms are located between the extremes and the middle of the segment.

However, the assumption of exogenous prices does not allow us to deal with the usual intuition: a firm can improve its profit by differentiating its product from that of its competitors with the intention of cooling down price competition. Recently, Barreda *et al.* (2011) have implemented experimental spatial markets with *endogenous price-setting*. This work presents two interesting results:

- There is a positive relation between differentiation and price.
- Differentiation by location tends to be low.

Figure 2.3 presents the aggregate results of the location – and price-setting stages, respectively. The results appear to be robust to variations in the experimental conditions regarding the rule of division of demand in case of a tie (automatized *vs.* human consumers).

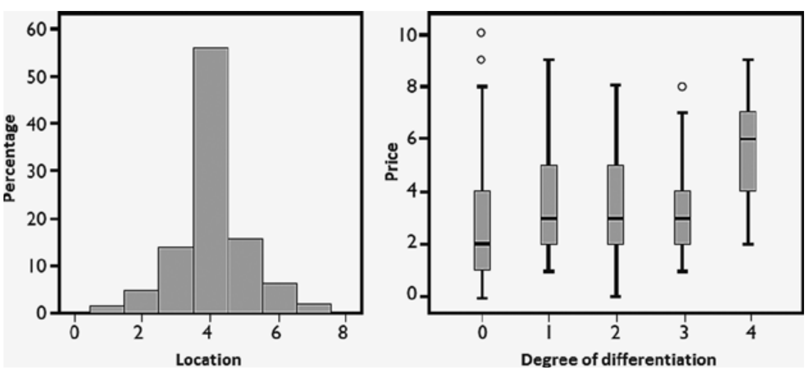

Figure 2.3 Spatial location with endogenous prices

Note: Left: The prediction of minimal differentiation by Hotelling (1929) is confirmed by the modal choice in an experiment with endogenous prices. Right: The hypothesis that prices increase for higher levels of differentiation is confirmed.

Source: Barreda *et al.* (2011)

In summary, the prediction of Hotelling's model, arguing that firms will agglomerate in the center of the market, is confirmed: both in experiments, where it is possible to compete in prices, as well as in contexts where the price is exogenously fixed. In addition, when price competition is allowed, subjects choose a product-differentiation strategy: they recognize its ability to sustain higher prices. It is very satisfying for a researcher to observe the predictive power of a theory in such a complex setting as the two-stage game – location and price – in Barreda *et al.* (2011), where the predictions of Hotelling (1929) are confirmed.

Finally, Camacho-Cuena *et al.* (2005) move one step further: spatial markets with endogenous prices, but in a context where the *location of the consumers is also endogenous*. Their experiments show that trial-and-error learning is not enough for consumers to understand that they should locate in the middle – between two sellers – which would increase competition and reduce prices.

In both experiments with endogenous prices mentioned here, coordination failures constitute an obstacle for sellers who attempt to avoid central locations to differentiate from other sellers.

Oligopolies with fixed quantities and extensions

The oldest oligopoly theory is the model of Cournot (1838), which differs from what has been surveyed up to this point due to the assumption that *firms compete by quantity* rather than price. The markets of Cournot (1838) have been experimentally studied by Fouraker and Siegel (1963), Holt (1995), Rassenti *et al.* (2000), and Huck *et al.* (2000, 2001b).

As discussed at the beginning of the chapter, a very general result of Cournot's experiments is that subjects, by adopting trial-and-error learning, show behavior that tends to confirm the theoretical predictions of Cournot-Nash. Although this is true on average, there is also certain variability around these predictions. These persistent oscillations decrease the predictive power of the equilibrium. In addition, in repeated-game scenarios, the total quantity is frequently not significantly different from the collusive prediction. In some cases, the total quantity oscillates between the collusive quantity and that of Cournot.

In summary, although moderately positive results regarding the predictive power of Cournot equilibria have been found, it seems that there are fundamental differences in the patterns of the data obtained in the experiments based on *Cournot models* (strategic quantity-setting) and the ones based on *Bertrand models* (strategic price-setting).

Huck *et al.* (2000) and Altavilla *et al.* (2006) both study markets of price and quantity-setting in a framework of differentiated products, where the Bertrand equilibrium is quite close to the Cournot equilibrium. Even though such experiments generally provide evidence that supports the Nash equilibrium predictions for the two types of markets (price-setting and choice of quantity), these works show that information derived from past strategies and results plays a crucial role.

Collusion in Cournot markets, *vs.* Bertrand markets, deserves thorough discussion. Suetens and Potters (2007) show that behavioral outcomes in Cournot markets tend to be more competitive relative to equilibrium as compared to those in Bertrand markets. Hence, more collusive behavior is detected on average in price-setting than in quantity-setting oligopolistic markets, with prices in price-setting experiments being above equilibrium prices, and quantities in quantity-setting experiments being above equilibrium quantities (see also Holt 1995, and Engel 2007). Moreover, the scope for tacit collusion in both types of markets is strongly affected by the number of competitors. Basically, implicit collusion is frequently observed in markets with two firms, rarely in markets with three, and almost never in markets with four or more firms. This effect has been observed under both Cournot competition (see, e.g., Huck *et al.*, 2004b) and Bertrand competition (see Abbink and Brandts, 2005, 2008; Orzen, 2008).

A famous modification of the Cournot model of quantity competition is due to Stackelberg (1934), which assumes that one of the two firms (the leader) chooses and commits in a credible way to produce a certain quantity, before the second firm (the follower) chooses its quantity. This model predicts that total production will be greater than that of the symmetric Cournot model and that, in general, the leader will produce and earn more than the follower. Huck *et al.* (2001a) experimentally compare the markets of *Cournot* and *Stackelberg*. Their results confirm the existence of volatility in the quantities chosen, similar to what is obtained in other experiments of production decision-making. In addition, the prediction that Stackelberg markets produce a higher total quantity is confirmed and that, therefore, such markets are more efficient than Cournot markets. However, the asymmetric nature of the interaction between the leader and the follower seems to hinder the convergence of the observed behavior towards the theoretical predictions of the Stackelberg model.[8]

Other experiments allow subjects to endogenously choose the moment they would make their strategic decision.[9] That is, before choosing quantities, the firms choose when to produce: in an early or a late period. In these experiments, simultaneous (Cournot) and sequential (Stackelberg) oligopolies can emerge endogenously. On the other hand, Huck *et al.* (2006) experimentally study a spatial market with *endogenous time choice*. In line with actual political campaigns, candidates can decide endogenously when and where to locate. Their results show that allowing for endogenous timing can eliminate some of the more unappealing equilibrium characteristics of the standard model.

In all these environments, the typical asymmetric results corresponding to the leader-follower structures receive less support than expected. In fact, the evidence is in favor of the symmetric results. This even occurs when the corresponding equilibria predict structures of the leader–follower type.

Cournot markets with *multiproduct firms* have received less attention in the experimental literature than corresponding Bertrand markets. A recent experiment by Hinloopen *et al.* (2014) analyzes the impact of product bundling in quantity-setting oligopolistic markets. One firm has monopoly power in a first market but competes with another firm à la Cournot in a second market. They compare treatments where the multi-product firm: always bundles, never bundles, and chooses whether to bundle or not. They also contrast simultaneous (Cournot) to sequential (Stackelberg) moves in the duopoly market. Their data support the theory of product bundling: with bundling and simultaneous moves, the multi-product firm offers the theoretically predicted number of units. In the case of sequential moves, when the multi-product firm is the Stackelberg leader, the predicted equilibrium is better attained with bundling, although this equilibrium is the same with and without bundling.

Oligopolistic models with *vertical relationships*, such as those between a manufacturer and a retailer, have received less attention by experimentalists than models with horizontal relationships. The basic model predicts that firms that are vertically related and not coordinated will set a total price–cost margin higher than if they agreed on a joint margin. This phenomenon is known as *double marginalization*. Durham (2000) experimentally confirms the importance of the double marginalization phenomenon predicted by the theory. Martin *et al.* (2001) corroborate these ideas in the lab and confirm the predictions of the theory of vertical integration as a means to exclude competitors. Furthermore, Normann (2011) shows that Bertrand markets with a vertically integrated firm are significantly less competitive than those where firms are separate. However, he shows that, while his experimental results violate the standard Nash-equilibrium notion, they are consistent with the quantal-response generalization of Nash equilibrium.

Other interesting extensions of the basic oligopolistic model are concerned with the *value of delegating* firm's strategic decisions to managers whose incentives are designed to either pursue or deviate from the firm's profit maximization. For instance, rewarding a manager according to the firm's production, rather than according to its profits, makes the firm a Stackelberg leader. Huck *et al.* (2004a) test the influential theory of Vickers (1985) and that of Fershtman and Judd (1987) on the strategic role of delegation in oligopoly through the design of incentives for the managers. These experiments study a situation in which the owners of the firm choose whether their managers receive as compensation either a share of the firm's profits or a reward based on the firm's revenues. Surprisingly, the second option is chosen very few (5%) times.

Georgantzís *et al.* (2008) study an *endogenous compensation system* in an experiment where the firm owners can choose between offering their managers compensation dependent on a linear combination of their own firm's profits and revenues, or compensation dependent on their own firm's profits relative to the profits of rival firms. As the theory presented in the paper predicts, the preferred remuneration system is the one based on own-firm profits, compared to those of rival firms.

Other (ir)regularities

A recurrent aspect in all oligopoly experiments, although it does not constitute a central topic in most of them, is *learning*. In fact, the trial-and-error learning process determines, to a great extent, which outcome is obtained in an oligopolistic market.

Cyert and DeGroot (1973) made an important contribution in this direction, by relating learning to the ability duopolists have in reaching a collusive outcome. It may be that learning in this context does not imply a disclosure of the mathematical properties of the supply and demand model the firm faces. In simple terms: learning is a dynamic process in an individual's decision-making. It is a process that leads him/her from initially uninformed strategies towards the relevant region of cooperative or non-cooperative equilibrium.

Multiple studies have attempted to identify possible systematic patterns in the learning strategies people have. Many researchers have analyzed adaptive learning (see, e.g., Nagel and Vriend 1999). A common result is that it does not seem that people learn through sophisticated or formal processes. For instance, García-Gallego (1998) and García-Gallego and Georgantzís (2001a) offered participants the opportunity of obtaining linear estimates (from an ordinary least squares model) of the underlying demand and, also, gave them various graphical representations (quantity-price, profit-price, etc.) of the data from the previous market period. The conclusion in both studies is clear: subjects did not make any effort to systematically calculate the optimal strategy by using explicit optimization, despite the fact that participants were academically advanced students (some of them were even graduate students in economics).

However, learning significantly affects observed behavior, as in most experiments participants' strategies first show a high degree of dispersion (they look almost random), and evolve over time towards the reference solutions, such as the collusive or non-cooperative outcomes. This brings us to an important determinant of the collusive outcome in oligopoly experiments. Mason and Phillips (1997) confirm the importance of information in a duopoly with asymmetric costs. In general, it is important to distinguish between *two sources of information* provided to participants in an experiment:

- Information can be provided *ex-ante*, through the instructions (at the beginning of the experiment).
- Information can be provided *ex-post*: that is, it becomes available as a result of past choices (feedback given during the experiment).

Contrary to what has been suggested by theorists, experimental treatments where information is given prior to the experiment have very little or no effect at all upon observed behavior: subjects hardly ever use information that is not immediately interpretable in their decision-making process. Therefore, information about the exact conditions of supply and demand, for instance, has very little effect on observed behavior. This is especially true when these conditions do not provide a "linear" interpretation to supply and demand.

On the contrary, informing participants about the strategies chosen, and even more importantly, about the outcomes obtained by their competitors in previous periods, has a significant effect.

Knowing the outcome of previous choices is not the only source of learning in oligopoly experiments. Various experiments have identified learning processes different from those assumed in the theoretical models. For instance, Huck *et al.* (1999) show that information about other players' strategies plays an important role in the emergence of collusive (rather than non-cooperative) outcomes. In addition, imitation of the most successful competitors appears to be supported by some experimental evidence.[10]

Another important aspect, that systematically affects behavior in oligopoly experiments, is *inequity aversion* (see chapter 6 of Vol. 1). Inequity-averse subjects tend to pursue payoffs similar to those of the competitors, even when the experimental context is initially asymmetric. Admitting a certain level of participants' inequity aversion can explain all those cases in which the theory fails to predict the observed behavior in asymmetric oligopoly experiments. The result of Huck *et al.* (2001b) explicitly relates inequality in payoffs with lack of stability in the Cournot setting. Altavilla *et al.* (2006) find that informing the oligopolists about past prices set by their competitors leads to quantities closer to those predicted by the Cournot-Nash equilibrium, while providing information about the average profit of the entire industry leads to higher levels of cooperation.

Aspiration levels of oligopolists are also constitute an important factor when looking at the divergence between theoretical and experimental outcomes. Huck *et al.* (2007) show that aspiration levels can be used to explain the merging paradox – where the merged firm ends up earning less despite there is less competition in the market after the merger – observed in the laboratory. Indeed, the aspiration-level hypothesis predicts that after the merger a firm has a target profit in mind, the one obtained before the merger, and it will also act in order to maintain it after the merger.

Finally, the role that *risk aversion* has on strategic behavior is evident in Sabater-Grande and Georgantzís (2002), showing how more risk-averse individuals have a lower probability of cooperating in a Prisoner's Dilemma (see chapter 7 of Vol. 1). The latter can be interpreted as a limit case (with only two strategies per player) of a Bertrand or a Cournot standard duopoly. The management literature is full of business cases identifying over-risky decisions made by managers. If we acknowledge that the managers may deviate from the pure profit maximization due to their aspiration levels, for personal and psychological reasons, or due to the incentives in their management contracts, the idiosyncratic effects observed in oligopoly experiments can be especially relevant to decisions that firms make in the real world.

The issue of whether the results of laboratory experiments on oligopolistic markets also extend to comparable (real) situations *outside the laboratory* is discussed in Potters and Suetens (2013). They correctly emphasize that there are several dimensions to this concern. In particular, decision-makers in firms are not students, and firms' decisions are usually not made by one individual acting on his/her own behalf. As to the latter concern, some experimental studies have begun to explore decisions by groups of individuals (boards) and how these depend on the decision-making process in the group. As Potters and Suetens (2013) report, in some cases individuals seem to act very differently from groups, whereas in other cases few differences are found.

Conclusions

Experiments on oligopolistic markets have aimed at testing the predictive power of oligopoly theory in explaining observed behavior in experimental settings that implement the conditions established by each model. Although such aims may appear to have limited relevance to the world outside the laboratory, this line of research has taught us some very interesting behavioral principles. For example, the oligopolistic equilibrium may be the limit towards which the strategies of economic agents who learn by trial and error converge. Further, such a learning process has a higher probability of supporting symmetric than asymmetric theoretical predictions. Finally, this learning process, as well as some idiosyncratic features of the experimental participants, can either help or hinder the corroboration of the theoretical predictions. Throughout this chapter we have tried to make it clear that we remain closer to the beginning than the end of this thrilling process of understanding how imperfectly competitive markets work.

Notes

1. According to the Greek etymology of the word, oligopoly (*ολιγοπωλιο*) is a market with few sellers (*ολιγοιπωλητές*).
2. The first experimental tests of Industrial Organization Theory can be found in Plott (1982). Later, Holt (1995) presented a summary of the experimental results in oligopolistic markets.
3. Kelly (1995) is one of the first examples of a multi-product monopoly experiment.
4. See for instance García-Gallego *et al.* (2004).
5. In other experiments subjects are provided with discrete payoff matrixes, which are the reduced version of the original oligopoly games with continuous strategies. The choices for a more or less realistic experimental design depend on the principles followed by the experimentalist and on his/her research objectives.
6. Fonseca and Normann (2008) provide a thorough analysis of the impact of *mergers* in experimental Bertrand oligopolies. They consider as treatment variables the number of firms (two, three) and the distribution of industry capacity (symmetric, asymmetric). They find that, even though they are more concentrated, asymmetric markets exhibit lower prices than symmetric markets with the same number of firms. Consistent with the static Nash-equilibrium prediction, duopolies charge higher prices than triopolies. However, although the overall impact of a merger is anti-competitive, the price increase is not significant. This last result, in a sense, confirms the findings of Davis and Wilson (2005).
7. Price parallelism in uniproduct firms has been previously studied by Harstad *et al.* (1998) in a context specifically designed to tackle this question. It was found that the conscious adoption of price parallelism by the competitive sellers had the effect of increasing prices towards the collusive prediction.
8. Kübler and Müller (2002) analyze experimentally markets with price-setting designed with the aim of comparing simultaneous and sequential decisions. They highlight the difference between authentic sequential games and sequential strategies obtained through the "strategy method" (subjects are asked what they would do for each choice profile of the other subjects and then the binding case is randomly chosen).
9. Huck *et al.* (2002a), Fonseca *et al.* (2005, 2006), and Muller (2006) are some of the experimental studies; see also Normann (2002) for a theoretical analysis.
10. See Offerman and Sonnemans (1998), Offerman *et al.* (2002). For a contrary option see Bosch-Domènech and Vriend (2003).