



UvA-DARE (Digital Academic Repository)

Bounded rationality and learning in market competition

Kopányi, D.

[Link to publication](#)

Citation for published version (APA):

Kopányi, D. (2015). *Bounded rationality and learning in market competition*. Amsterdam: Tinbergen Institute.

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

Chapter 1

Introduction

The main purpose of economic models is to provide a framework in which the effect of different factors on the economic outcome can be analyzed. Will competition lead to an efficient allocation? Should the government intervene in the market in order to protect consumers? Which policy instruments should be used to reach the desired outcome? In order to get adequate answers to such questions, it is important to understand how actual markets work and how to model markets properly.

Mainstream economic models are built upon the assumption of rationality. This means that agents perfectly know the environment in which they interact with each other (or they have correct beliefs about it), they have correct beliefs about the actions of other agents and they have the cognitive and computational abilities to determine the optimal action given their preferences and beliefs.¹ Under these assumptions, agents do not make systematic errors and the rational choice is the predicted outcome. In game theory, standard equilibrium concepts such as Nash equilibrium, Bayesian Nash equilibrium and sequential equilibrium rely heavily on the rationality assumption. Players can calculate their best response function and they have correct (or consistent beliefs) about the actions of the other players. The same holds true for the field of industrial organization: firms as well as consumers are assumed to be rational.

¹Note that rationality does not impose restrictions on preferences. Thus, we do not consider behavior as irrational if agents pursue alternative objectives than simply maximizing their own payoff.

An often used argument in favor of the rationality assumption is that agents *learn* to behave rationally. Even if agents do not know what the optimal (or equilibrium) action is, reasonable learning methods will lead to the equilibrium, therefore the rationality assumption can be applied. However, there are many different learning methods and different methods may lead to different outcomes. Therefore the learning argument is not necessarily valid for supporting the rationality assumption. Another argument, due to Friedman (1953), is that irrational agents will be driven out of the market by rational agents. This argument sounds quite reasonable: irrational agents make bad choices, therefore they will lose their wealth. Note, however, that there are different degrees of (ir)rationality therefore comparing fully rational and fully irrational agents might not be the most relevant case. The bounded rationality approach assumes that even though agents are not perfectly rational, they make a reasonable choice subject to their abilities, information and possibilities. Friedman's argument is not necessarily true if we compare perfectly rational and boundedly rational agents. In fact, Schaffer (1989) shows that it is typically not the profit maximizers that have the highest survival rate in a natural selection type framework that Friedman used in his argument. This is explained by spiteful behavior: by deviating from profit maximization, a firm can increase its chance of survival by choosing an action that hurts itself but hurts its competitors even more. Schipper (2009) gives another example where firms that simply imitate the behavior of the best-performing firm make a larger profit in equilibrium than the more sophisticated optimizers. Moreover, Haltiwanger and Waldman (1985) illustrate that boundedly rational agents can affect the equilibrium outcome more than proportionally to their weight in the population. This emphasizes that bounded rationality must be taken into account.²

The assumption of rationality can be quite demanding informationally as well as computationally. In a market model, consumers should know which goods are supplied on the market, they need to know all the relevant characteristics of these goods (price and quality for example) and they must be able to choose the combination of goods that maximizes their utility.³ Simi-

²Conlisk (1996) overviews and criticizes other typical arguments in favor of perfect rationality.

³Of course, there is a tradeoff between informational and computational burdens. There are models where consumers are not assumed to know all the prices or qualities of goods and they need to search to gather these pieces of information. Finding the Nash equilibrium in these models is much more demanding computationally as

larly, firms should perfectly know how the demand for their good depends on the price or production choice of their competitors and they should have correct beliefs about their competitors' choice. These assumptions are quite strong. Agents might not know the market environment perfectly, they might not understand the market mechanism correctly and they might not have the computational abilities to determine the optimal action.

There is ample empirical evidence showing that agents' behavior is not necessarily in line with predictions under the rationality assumption. Conlisk (1996) provides a good overview of such results. Armstrong and Huck (2010) survey further experimental evidence as well as real-world examples supporting the relevance of bounded rationality. Since models of bounded rationality may describe actual behavior better than perfect rationality in some situations, bounded rationality should be taken into account in theoretical market models. Therefore in this thesis we consider market models where firms are boundedly rational and we investigate the effects of bounded rationality on the market outcome.

Bounded rationality

Bounded rationality concerns modeling behavior that deviates from perfect rationality. In models under perfect rationality, agents' decisions are derived from a maximization problem such as utility or profit maximization. In complex environments this task can become very complicated and agents might not have the cognitive or computational abilities to solve their optimization problem. In his survey paper, Ellison (2006) distinguishes three approaches to bounded rationality. The first one is a rule-of-thumb approach in which agents follow simple and reasonable rules for decision making instead of solving optimization problems. In the second approach agents explicitly take cognitive costs into account in their optimization problem, whereas the third approach incorporates behavioral biases found in the psychological and economic literatures.⁴

consumers need to calculate the price or quality distribution in equilibrium to determine the optimal decision rule.

⁴Strictly speaking, we do not consider the second approach as a form of bounded rationality. If there are bounds on rationality then solving an even more complicated optimization problem does not seem a good modeling approach. Selten (1990) also argues that bounded rationality is not related to an optimization problem.

The use of bounded rationality was first suggested by Herbert Simon. Simon (1955, 1972) points out that agents might have limited computational abilities or limited information gathering capacities, therefore they may not be able to find the optimal choice. In this situation agents might be *satisficing* instead of optimizing: they use a simple procedure to find an outcome that yields a payoff that they consider satisfactory. Tversky and Kahneman (1974) provide empirical evidence that simple decision rules can describe subjects' behavior well. Of course, there may be many different decision rules that are reasonable in a given situation and agents may prefer to use different rules. In fact, Stahl (1996) finds experimental evidence both for heterogeneity among individuals and for switching to rules that performed better in the past. Moreover, agents may update the rules that they are using.

Bounded rationality has been applied in many branches of economics. Conlisk (1996) provides a good overview about the different economic applications of bounded rationality. Ellison (2006) focuses on the use of bounded rationality in industrial organization. Even though the early models concerned bounded rationality on the firm side, most models focus on boundedly rational behavior on the consumer side while firms are assumed to be rational. Spiegel (2011) also discusses several models where consumers are boundedly rational. In contrast, Armstrong and Huck (2010) give an overview of the literature on the behavioral approach on the firms' side. Examples include satisficing behavior, imitation and alternative objectives that deviate from profit maximization. This thesis also contributes to the literature on bounded rationality on the firms' side. Kirman and Vriend (2000) provide a model of a fish market where both the consumers and the sellers are boundedly rational. Their model captures important characteristics observed in the fish market in Marseille: price dispersion and the loyalty of buyers to sellers.

We can conclude from the literature that bounded rationality is relevant in some situations but less so in others. Even though this conclusion is not strong on its own, it clearly shows that bounded rationality should be taken into account as it is capable of describing actual behavior. In fact, it may describe actual behavior better than a model under perfect rationality.

Learning

Besides bounded rationality, another factor that can have an important effect on the market outcome is learning. Learning can be incorporated in models in different ways. Agents can learn about the market environment, about other agents' behavior or about the optimal action.⁵ These categories are not exclusive, all three can be incorporated in the same model.⁶ Consider for example a dynamic market model where firms do not know the demand function for their product. With a certain learning method they can learn the demand conditions from the market history (using observations about prices or production levels). Once they learned the demand conditions, they might use observations about their competitors' action to learn (predict) their action for the next period. When firms know the demand conditions and they have a prediction about their competitors' action, they need to determine the optimal action subject to their beliefs. If they cannot calculate the best response for some reason, then they may apply a learning method to find out what the optimal choice is.

There exists a wide variety of methods for modeling learning in the economics literature, including different belief-based models, adaptive learning and evolutionary methods. Fudenberg and Levine (1998), Evans and Honkapohja (2001) and Cressman (2003) give good overviews of these classes, respectively. Different learning methods may lead to different outcomes. This is illustrated in Offerman et al. (2002), for example: they consider two imitation-based and one belief-based learning rule that lead to different market outcomes theoretically as well as in a laboratory experiment. This shows that it is essential to explicitly model the agents' learning behavior. Furthermore, the heterogeneity of agents should also be taken into account, as promoted by Hommes (2006, 2013) for example. Agents may prefer to use different learning methods (due to differences in computational abilities, for example) for finding out what the optimal decision is. Therefore, it is important to analyze what happens in a heterogeneous

⁵Note that the last two learning categories do not exist under perfect rationality. When each agent is rational, agents do not need to learn each other's behavior as they know it by default. Also, agents can calculate the optimal action so there is nothing to learn about it either.

⁶They are not exclusive either in the sense that the same learning method can be used to learn about different objects as well.

environment and how different learning methods affect each other.

In this thesis we focus on learning about the market environment. There are many situations where firms do not have full information about the environment in which they operate. For example, they might not know how the demand for their good depends on the price they charge, how it is affected by their competitors, who their competitors are and how they act. Learning is especially important when the market is subject to a structural change (e.g. a new product is introduced or a new firm enters the market). Learning has a natural role in these situations: firms gather the information resulting from their actions, they evaluate it and take it into account when making a decision. A natural way to learn about demand conditions is to gather market data and use it to estimate a demand function. This can be modeled with *least squares learning* (LSL). LSL consists of two parts: estimation and a decision rule. In a given period, firms use their past observations about prices and demands to estimate the unknown parameters of a so-called *perceived demand function*. Then, based on the parameter estimates, they choose the action that maximizes their perceived profit. When a new observation arrives, firms update their parameter estimates, leading to (possibly) different prices.

Thesis outline

In this thesis we incorporate bounded rationality and learning on the *supply* side of the market and we analyze how this affects the market outcome. We do not deviate from rationality on the demand side and we summarize the consumers' behavior by a demand or inverse demand function. We focus on two aspects of rationality in the analysis. First, we relax the assumption of perfectly knowing the market environment. In Chapters 2 and 3 we consider different learning methods firms may apply and we analyze the corresponding market outcome. Second, we investigate the effects of deviating from rationality through the beliefs about the actions of other firms. We consider situations where firms perfectly know the market environment but they are uncertain about their competitors' actions. In Chapter 4 we analyze the consequences of relaxing the consistency requirement on beliefs about competitors' actions while in Chapter 5 we

investigate how the market outcome depends on the type of information agents receive about the other agents' actions. We apply analytical as well as numerical methods in the thesis. We analyze the models theoretically as much as possible but we have to use computer simulations and numerical methods when formal analysis is not tractable. Moreover, we analyze behavior by means of a laboratory experiment in Chapter 5.

In Chapter 2 we focus on the interaction between different learning methods. We consider a Bertrand oligopoly with heterogeneous goods where firms do not know the demand function and they can apply two learning methods to determine their price. One of the methods is a misspecified version of least squares learning, with which firms focus on their own price effect only, thus they do not take into account all the relevant variables that affect the demand for their good. Under this learning method firms typically reach a so-called *self-sustaining equilibrium*. This term is introduced by Brousseau and Kirman (1992) and it denotes a situation in which firms maximize their profit subject to their beliefs about the demand conditions and these beliefs are correct at the equilibrium point but not outside the equilibrium. That is, firms do not learn the true demand function correctly. The other learning method is gradient learning: firms adjust their price in the direction that gives a higher profit. This method leads to the Nash equilibrium when it converges. In the steady states of the model where both learning rules are present in the market, least squares learners are in a self-sustaining equilibrium while each gradient learner gives the best response to the prices of other firms. When firms are allowed to switch between learning rules based upon how well the rules perform, we can observe a cyclical switching that is driven by changes in the stability of gradient learning.

In Chapter 3 we further investigate the properties of misspecified least squares learning. We consider the circular road model of Salop (1979), with three firms and two types of consumers. One consumer group faces low transportation costs while the other group has high transportation costs. Firms do not know the demand structure of the model and they apply least squares learning to learn the demand conditions. In the estimation the firms use a perceived demand function that linearly depends on the prices. In contrast to Chapter 2, firms can observe the

prices of each other and they use this information in the regression. Since the true demand function is piecewise linear, perceived demand functions are correctly specified locally but not globally: firms can learn at most one linear part correctly. We prove that the model has three kinds of equilibria. Firms may reach a self-sustaining equilibrium, the Nash equilibrium or an asymmetric learning-equilibrium in which one firm focuses only on the consumers with high transportation costs whereas the other two firms serve both consumer types. Both the Nash equilibrium and the asymmetric learning-equilibrium are locally stable, therefore the model has coexisting stable equilibria. We analyze the conditions under which the different outcomes are reached.

In Chapter 4 we investigate the effects of having weakly consistent beliefs about the competitor's action. We consider a duopoly where firms produce a homogeneous good and they set *both* the price and the production level of the good simultaneously. In the standard model under perfect rationality, there exist a Nash equilibrium only in mixed strategies but not in pure strategies. In our model, firms are risk averse and they hold probabilistic conjectures about the actions of the other firm. We numerically show that our model may have an equilibrium in pure strategies. Beliefs are weakly consistent in this equilibrium: the *modes* of the belief distributions correspond to the actual actions of the other firm. We investigate how the degree of risk aversion and the amount of uncertainty regarding the price and production level of the other firm affect the equilibrium. Our results show that a small degree of risk aversion is welfare enhancing and that welfare is higher than in the mixed-strategy equilibrium of the standard model. The chapter illustrates that having weakly consistent beliefs about other agents' actions might lead to a substantially different model prediction than under the assumption of consistent beliefs and that a small amount of bounded rationality may be welfare enhancing.

In Chapter 5 we investigate by means of a laboratory experiment how the type of information firms receive about past choices of their competitors affects their behavior and the market outcome. We conduct a laboratory experiment in which subjects play the role of firms in the market. We vary the amount and the type of information subjects receive about the choice

of their competitors. Subjects either receive or do not receive information about production levels. When subjects are informed about quantities, they can see either the total output or firm-specific production levels. Moreover, we introduce voluntary information sharing in one part of the treatments, where subjects can choose to give information about their past production level to their competitors. Our results show that subjects use information sharing to signal their willingness to collude. Aggregate outputs tend to be lower under individual information than under aggregate information, supporting the view of competition authorities that the publication of firm-specific information has anti-competitive effects. Under perfect rationality, the equilibrium prediction is the unique Nash equilibrium of the one-shot market game. However, we can observe substantial dispersion in the market outcomes in each treatment. Therefore this chapter provides experimental evidence that the rational prediction may not describe individual behavior well.

Finally, the main findings of the thesis are summarized in Chapter 6.