MAGISTERARBEIT

Titel der Magisterarbeit

„Agent-based models of Cournot oligopoly"

Verfasst von

Simon MARTIN

angestrebter akademischer Grad

Magister der Sozial- und Wirtschaftswissenschaften
(Mag.rer.soc.oec.)

Wien, 2014

Contents

1 Introduction

In the traditional economic analysis of firm behavior, three key quantities play a major role: The monopoly quantity Q^M, the Walrasian (or fully competitive) quantity Q^W, and the Cournot quantity Q^C. According to the theory, if a monopolist serves the market, he chooses to produce the quantity Q^M, at the expense of the consumers. On the other extreme, if the market is fully competitive, each firm chooses a quantity such that the price equals the firm's marginal cost, resulting in zero profits for all firms. Between those two extremes is the Cournot quantity Q^C, representing the total production quantity chosen by perfectly rational firms with perfect foresight in an oligopolistic market.

Huck et al. (2004b) make experiments with participants representing two, three, four and five firms, in order to analyze the competitiveness of Cournot markets. They find that in a setting with two subjects, there is a tendency towards collusion, i.e., producing a quantity below the Cournot output and thus reaping higher profits. For three or more subjects, they find a tendency towards the Cournot solution. In the experiments of Brandts and Guillén (2007), two subjects decide not only on the production quantities, but also on the price. They find that the subjects either end up in a pure monopoly (because the other ones are forced to leave the market due to unprofitability), or collude at monopolistic output levels.

Attempting to model firm behavior in a way that could explain why firms might in fact do better than what the standard Cournot model predicts, Huck et al. (2004a) make firms maximize absolute profits by following a rule called "trial and error" in a continuous time setup. If a certain direction of movement is beneficial for the firms, they simply keep that direction. Huck et al. find that firms learning according to "trial and error" converge to the joint-profit maximizing solution, i.e., the total production quantity equals exactly the quantity a monopolist would choose (i.e., the total production quantity $Q = Q^M$). An earlier version of the paper, Huck and Oechssler (2000), finds the same results, but in discrete time.

An agent-based model of Kimbrough and Murphy (2009), however, leads to the conclusion that firms behaving according to an algorithm called "probe and adjust" converge to the Cournot solution (i.e., the total production quantity $Q = Q^C$). "Probe and adjust" is similar to "trial and error", as firms maximize absolute profits and explore the production space around quantities which they found profitable already.

In the dynamic Cournot model of Vega-Redondo (1997), firms do not seek to maximize absolute profits, but relative profits (i.e., market share) instead. Vega-Redondo argues that this best reflects the firm's main objective of survival. Firms achieve relative profit maximization by imitating last period's most successful firm. For any finite number of firms, the total production quantity equals the Walrasian outcome (i.e., $Q = Q^W$).

In this paper, I reconcile the two models of Huck and Oechssler (2000) and Kimbrough and Murphy (2009) and explain the key differences in their models, resulting in the contradictory results. I show that rigidity of "trial and error" firms allows them to do better than "probe and adjust" firms, who do not optimize in such a straightforward way, but rather probe certain production quantities and attempt to improve from there on. Once the assumption of pure rigidity is relaxed and "trial and error" firms start experimenting randomly, also they converge to the Cournot outcome.

Firms in the model of Vega-Redondo require perfect knowledge for their imitation behavior. In order to analyze whether also firms with imperfect knowledge converge to the Walrasian outcome of Vega-Redondo, I modify the "probe and adjust" algorithm by making firms seek to maximize relative profits rather than absolute profits. The results found show that also "probe and adjust" firms, maximizing relative instead of absolute profits, converge to the Walrasian outcome.

The remainder of this paper is structured as follows. In Section 2, I give an overview about previous literature on the subject. First addressing the question of the contradictory results of Kimbrough and Murphy and Huck and Oechssler, Section 3 explains the two models in detail. Section 4 describes the replication of the two models. In Section 5, I gradually modify the assumptions of Kimbrough and Murphy towards the ones of Huck and Oechssler until I find matching results. An analysis of the sensitivity of "trial and error" firms to the propensity to experiment is performed in Section 6. In Section 7, I make "probe and adjust" firms optimize relative profits and compare the outcome to the results of Vega-Redondo. Section 8 concludes.

2 Literature review

Traditional economic theory typically employs three major tools to analyze markets with oligopolistic competition: the Cournot model, the Bertrand model, and the Stackelberg model (Gravelle and Rees, 2004). All of them treat the respective decisions of the firms as one-shot games. Firms are assumed to be perfectly rational and to have perfect foresight.

In the Cournot model, firms simultaneously decide on the quantities to produce. Once their binding decision is made, prices are determined according to market demand, and firms realize their respective profits. Assuming perfectly rational competitors, firms end up in the Nash equilibrium. Total production quantities are higher than what a monopolist would choose, but lower than the quantities of a perfectly competitive market, where prices equal marginal costs (Gravelle and Rees, 2004).

Conversely, in the Bertrand model, firms do not make quantity decisions, but price decisions instead. The quantities produced and sold are determined by market demand upon binding simultaneous price decisions of the competitors. In a monopoly, it does not matter whether a price or a quantity decision was made; in an oligopoly, however, there is a difference. For homogeneous products, the Bertrand outcome are prices equal to the marginal costs, i.e., the competitive market outcome, sometimes also referred to as the Walrasian result (Gravelle and Rees, 2004).

In his seminal work, Vega-Redondo (1997) analyzes the evolutionary process of a dynamic Cournot model with finitely many firms competing in a market with a homogeneous product. Contrary to traditional theory, he does not assume that firms seek to maximize absolute profits, but relative profits instead, i.e., payoff differentials. His justification is that in a Darwinian sense, survival is the primary objective of the firms. A competitor which gets stronger and stronger might eventually drive a firm out of the market, even though it made big profits on its own. Thus, we might view firm behavior as "responding to forces of learning and imitation" (Vega-Redondo, 1997, p. 382). It might result in "spiteful behavior", i.e., firms might be willing to take actions which harm themselves, as long as these actions harm their competitors even more. In Vega-Redondo's (1997) model, firms choose their quantities from a finite set in a discrete time setup. In each period, firms "imitate" the behavior of last period's most successful firm, or with a small

probability choose to experiment with any other randomly chosen quantity. Thus, the model can be seen as a model of dynamic learning. As in the standard Cournot model, firms take prices as given. However, in stark contrast to the outcome of the standard Cournot model, for any finite number of firms, in particular for a small number such as two or three, Vega-Redondo finds convergence to the Walrasian, i.e., fully competitive outcome, instead of the Nash equilibrium.

Also Huck et al. (2004a) investigate a dynamic Cournot model, but show a result on the other extreme, namely convergence to the monopoly solution. They model their agents to follow a learning process called "trial-and-error". The agents have only two options: increase or decrease the output quantity by a fixed amount. Afterwards, the agents assess whether the outcome was beneficial for them or not. Depending on the outcome, they continue their choices in the same direction or revert it. There are two important differences to the setup of Vega-Redondo (1997): (i) "trial-and-error" learning does not require any information about the other agents' actions, contrary to the agents in Vega-Redondo, who require full knowledge in order to imitate their opponents, and (ii) while in Vega-Redondo (1997) the success measure for agents were relative profits, Huck et al. make their agents optimize absolute profits. Huck et al. stress that agents' collusion is only possible because coincidentally, they find that simultaneous downwards movements from the Cournot solution are beneficial for all of them. The standard Cournot model does not allow for such random encounters. Huck et al. mention that already 1964, Baumol and Quandt obtained a similar result for a monopoly. While Huck et al. (2004a) work with continuous time, Huck and Oechssler (2000) provide an analysis in discrete time, finding the same results, i.e., that firms behaving according to "trial and error" learn to produce the joint-profit maximizing quantities.

Several authors analyze Cournot oligopolies in agent-based models, e.g., Waltman and Kaymak (2008), Arifovic (1994), Kirman and Vriend (2001) and Kimbrough and Murphy (2009). Using agent-based models allows capturing complex process dynamics that are out of scope of analytical frameworks. Kimbrough and Murphy attempt to connect economic theory, recommendations from management literature and the results of the experiments by setting up an agent-based economic model using the modeling software NetLogo[1]. They want to analyze "the question whether there are effective procedures, using realistically available information, that may actually be used by managers in oligopoly settings and that produce Cournot-improving outcomes" (Kimbrough and Murphy, 2009, p. 49) by modeling agents in an oligopoly to behave as recommended by Nagle and Holden (2006). Agents are set up to evaluate their output decision in each period against a predefined success measure, e.g., the own profit, the overall profit of the entire market, or a convex combination of those two. If the decision to increase output proved to be advantageous according to the success measure, the agent continues exploring further possibilities from there on. Thus, they call this process "Probe and Adjust". There are apparent similarities to the "trial-and-error" learning proposed by Huck et al. (2004a). The "Probe and Adjust" process is repeated until the agents' output decisions stabilize. If the agents' success measure are only (absolute) own profits, the model yields the Cournot solution, contradicting the findings of Huck et al. (2004a), who find convergence to the monopoly solution. It is not apparent immediately which differences in the models cause these contrasting results.

However, if agents follow a policy called "Market Returns, Constrained by Own Returns", a firm "looks to get its share of the market and then looks to keep the market as profitable as

---

[1]http://ccl.northwestern.edu/netlogo/

possible" (Kimbrough and Murphy, 2009, p. 52), and the model admits the monopoly solution.

This completes the description of the three papers which are most relevant for the analysis at hand. More generally speaking, the standard Cournot model has been criticized and extended in various ways in the literature.

Leonard and Nishimura's (1999) claim that in the Cournot model the assumption of knowledge of exact demand functions is unreasonable, because "knowledge is imprecise and information expensive" (Leonard and Nishimura, 1999, p. 165). Leonard and Nishimura analyze a duopoly model with a homogeneous nonstorable good and general nonlinear demand where they relax the aforementioned assumption by letting firms make error-prone estimates about the consumer's demand behavior. Both firms face linear total cost. As in the standard Cournot model, firms try to maximize profits by assuming that the competitor keeps his production quantity constant. Leonard and Nishimura find that when firms make mistakes in their estimates of consumer demand, these mistakes do not necessarily disappear, but might instead lead to a two-period equilibrium cycle.

Bischi et al. (2002) extend the model of Leonard and Nishimura by setting up a duopoly model with cost externalities among firms, building on the idea that the prices of an input factor might change for firm A if firm B demands more or less of this input factor. Again, agents do not possess perfect knowledge, but might have a biased picture of the world. Bischi et al. understand this imperfect knowledge as representation of the fact that firms might simply make errors during their estimates or have lagged information. The inverse demand curve is assumed to be linear, but the cost functions are nonlinear. Stressing that the results depend on the initial conditions, Bischi et al. find that "the economy ends up in subjective equilibria, because the outcome observed [by the firms] is perfectly consistent with their (mis-specified) beliefs" (Bischi et al., 2002, p. 20). The equilibrium established might even be "far away" from the Nash equilibrium, even if the mis-specification is relatively small, resulting in welfare loss.

Theocharis (1960) shows that there is a stable solution in a Cournot model with a linear inverse demand function and constant marginal cost when there are only two firms, but this solution becomes unstable once there are more than four firms. When there are exactly three firms, there are "finite oscillations about the equilibrium" (Theocharis, 1960, p. 133).

On the contrary, Rand (1978) finds effects which he refers to as "chaotic" even in the fully deterministic duopoly model in discrete time without production costs. This behavior emerges as a result of the construction of the reaction functions.

Furth (2009) states that, in general, equilibria need not exist; if they exist, they need not be stable; and if they are stable, they need not be unique. Quoting Debreu (1974) and Mantel (1974), he claims that "'everything is possible', sometimes formulated as: 'anything goes'" (p. 183). Furth shows that, even though "all kinds of instabilities" are possible in heterogeneous oligopolies, in homogeneous duopolies this need not be the case. He analytically proves that homogeneous duopolies admit these instabilities only if there are cost externalities, which Furth terms "unrealistic", or the reaction correspondences exhibit certain properties. However, he does not present an analytic proof for oligopolies with more than two firms.

In the setup of Vega-Redondo, the economy consisting only of "imitators" converges to the Walrasian equilibrium. The standard Cournot model consisting only of "optimizers", always playing the myopic best response to the last round's total output, converges to the Nash equilibrium.

Schipper (2009) mixes the two approaches by modeling an economy consisting both of optimizers and imitators in a repeated Cournot game. Schipper finds that in this model, imitators are generally "strictly better off than optimizers" (p. 1982). He offers the possible explanation that imitators are in a sense similar to the "first movers" in the Stackelberg game and have the according first-mover advantage, because they are (in a sense) "independent" of what the competitors do. Optimizers, on the other hand, are "dependent" on their opponents decisions in previous rounds. Schipper claims that this shows "ambiguous semantics of profit maximization" (p. 1987), as even though optimizers want to maximize their profits, their attitude is too myopic and therefore they cannot reap maximum profits. The question what would happen if optimizers were not myopic is left open.

The actual business environment which firms typically face is coined by non-perfectly rational producers and consumers and imperfect information. In particular, the exact demand behavior is rarely known to a firm. Management literature provides recommendations to firms operating in such an environment, e.g., Nagle and Holden (2006). Managers should "rather explore than optimize" (Kimbrough and Murphy 2009, p. 51).

Among the first to analyze a Cournot oligopoly in an agent-based simulation model was Arifovic in 1994 by employing a genetic algorithm (GA) which she summarizes as follows: "Rules whose application has been more successful in the past are more likely to become more frequently represented in the population, through a process similar to the natural selection in population genetics" (Arifovic, 1994, p. 4). She claims that the behavior generated by computer-based adaptive algorithms is more in line with data from experiments than models with rational models, i.e., the standard Cournot model, predicts, by converging to the fully competitive equilibrium. Even when the market price emerging in each period is not known, the agents find a solution as if the price were known. Arifovic's (1994) findings are mainly in line with Vega-Redondo (1997).

According to Riechmann (2002), Cournot models of evolutionary learning yield different outcomes, depending on the type of learning employed. If agents learn socially, i.e., they learn from each other, e.g., by interaction, the outcome is Walrasian, as predicted by Vega-Redondo (1997). Similar to Vega-Redondo, Riechmann attributes this to the "spite effect", i.e., the agents focus on relative payoffs instead of absolute payoffs and thus are willing to take actions which harm themselves, as long as they harm their opponents even more. Imitation of last periods most successful agent leads to the Walrasian outcome.

Conversely, if individual learning is employed, i.e., there is no interaction among agents, there is no imitation, but rather learning by "introspection", avoiding the harmful spite effect. In that case, according to Riechmann, the more rational agents are, the more likely they are to leave the Walrasian and converge to the Cournot outcome. However, as Riechmann refers to higher rationality if an agent is equipped with more knowledge, measured by "amount of information needed to conduct the retrospective learning method, the question if agents need a personal memory, and the computational abilities an agent needs (p. 18), the term "sophistication" seems more accurate than "rationality".

Waltman and Kaymak (2008) model repeated Cournot oligopoly games with Q-learning agents. Q-learning is a form of reinforcement learning, where agents decide on a finite set of actions based on the current state of the world. As in the model of Vega-Redondo (1997), agents have a certain probability to deviate and "experiment" instead. Watkins and Dayan (1992) proved that "... [Q-

learning] agents can learn to behave optimally under certain conditions" (Waltman and Kaymak, 2008, p. 3279). In the model of Waltman and Kaymak, there is a fixed number of firms, a linear demand function, identical marginal costs across firms, and the goods produced are perfect substitutes. Firms aim to maximize absolute profits. The repeated game discussed is different from the one-stage game, because firms maximize long-term profits, and thus "it may be possible to sustain collusion" (p. 3281). In their model, the probability of experimentation approaches zero over the time span of 1.000.000 periods. As all agents are Q-learners, no analytical result is possible. Waltman and Kaymak find a general tendency towards collusive behavior, but not full collusion.

The previously discussed models build upon the Cournot model; Zhang and Brorsen (2011) analyze an agent-based model for price-setting firms in an oligopoly, i.e., the Bertrand model. The analysis involved is much more complex, because small changes in prices can cause large changes in sales. Employing a method called Particle Swarm Optimization (PSO) and claiming that it finds better results than Genetic Algorithms (GA), their agents end up at the monopoly solution in a duopoly, and once more than four firms compete, the markets tend to become competitive. This is in line with the results from Huck et al. (2004b).

One of the most comprehensive agent-based computational economics (ACE) models to the day was created by Kirman and Vriend (2001), analyzing the wholesale fish market in Marseille by focusing on two stylized facts from a rich empirical data set: there is high loyalty of buyers to sellers, and there is persistent price dispersion. They model the market as a market hall with 10 initially identical sellers and 100 initially identical buyers, who meet in the market hall for 5000 days (periods). Buyers can choose which seller to ask for a price offer, and sellers can offer prices to each buyer individually. The buyer wants to buy one unit of fish every day, and he can decide to accept or reject price offers. As fish is a perishable commodity, inventories do not play any role in the analysis. Agents learn behavior by a reinforcement learning algorithm, i.e., a strategy that worked well in the past is more likely to be repeated in the future. Sellers are assumed not to have any reservation price. Loyalty of buyers to sellers and treatment of loyal customers are not embedded in the agents, but are shown to emerge.

Kirman and Vriend find that their model can explain the two aforementioned stylized facts very well. They interpret loyalty as a coordination device, being somehow similar to "intrinsically meaningless signals like wearing a blue shirt" (p. 491). However, a key difference is that loyalty emerges because agents find that it is beneficial for them, because, as they term it, "loyalty means continuity". The results are obtained even though agents are not forward looking and thus do not play dynamic strategies.

Barr and Saraceno (2005) analyze the effects of environmental and organizational effects on the outcome of repeated Cournot games with two firms by modeling firms as a network of agents, represented by a learning neural network which takes into account environmental factors. Firm sizes are exogenously given. They find that smaller firms are more flexible and therefore learn faster and perform better in the short run, whereas big firms perform better in the long run by finding more accurate solutions. The more complex the environment, the bigger the optimal firm size, because the short-run advantage of being smaller disappears faster. The firms represented by neural networks converge to the Nash equilibrium.

Kochanski (2009) develops a hybrid model containing both elements of agent-based simulation

and the analytical Cournot framework, analyzing determinants of market concentration. His model allows for marginal-cost reducing innovation and is consistent with what mainstream theory would predict - markets with high costs tend to be more concentrated. Kochanski (2009) demonstrates that analytical and agent-based model can complement each other.

The literature described above leaves the following questions open:

1. Kimbrough and Murphy (2009) and Huck and Oechssler (2000) find very contradictory results, even though the models used seem similar at first sight. How can one model find convergence to the Nash equilibrium, while the other finds convergence to the monopoly solution?

2. In the model of Vega-Redondo (1997), firms do not attempt to maximize absolute, but relative profits, i.e., market share. The same is true for the model of Riechmann (2002) when agents learn socially. Similar to Vega-Redondo, the firms converge to the Walrasian outcome when agents have full information. To complete our analysis, I will confirm that an economy with agents trying to maximize market share that do not have full knowledge also converges to the Walrasian outcome.

Both questions are addressed in the upcoming sections.

3   Model

In order to compare the contrasting outcomes of the models of Kimbrough and Murphy (2009) and Huck et al. (2004a), the following section describes the two models in detail. Afterwards, I set up my own model to match the described models as close as possible, and attempt to replicate the results found by the respective authors. I start with the model of Kimbrough and Murphy and use it as a reference, and afterwards introduce stepwise changes towards the model of Huck et al.. However, in 2004, Huck et al. analyze the model in continuous time, because "...[this] version lends itself to a much more elegant presentation and analysis" (Huck et al., 2004a, p. 207). Already in 2000, they discussed the model in discrete time, essentially finding the same results: "Theorem 1: For a duopoly the trial & error process converges to a collusive outcome. If cost functions are identical, then it converges to the joint profit maximizing outcome" (Huck and Oechssler, 2000, p. 7). As the simulation of Kimbrough and Murphy is, as most simulations, in discrete time by nature, further comparisons and development are much more straightforward to be done based on Huck and Oechssler (2000). Thus, I do not go into the details of Huck et al. (2004a), but only of Huck and Oechssler (2000).

As the strikingly different results appear even in the duopoly case, in the following I just describe and analyze the case of two firms.

3.1   Kimbrough and Murphy (2009)

This model is an agent-based simulation of a Cournot oligopoly. Even though the model provides several different success measures for the firms, in the following I only refer to the simulation where absolute profits of the firms were used as success measure (in terminology used by Kimbrough and

Murphy, the firms aim for "own returns"). As in the standard Cournot model, firms make quantity decisions, taking the price as given, which is determined by the market.

The agents have no information about the demand function or the other players decisions; they only observe the price which is determined by the market and consequently calculate their own profits. Demand follows a linear inverse demand function of the form

$$P_t(Q_t) = \max(0, \alpha - \beta Q_t) \tag{1}$$

where P_t is the market price in period t, α and β are constants and $Q_t := q_{1,t} + q_{2,t}$. In the following, the subscript t is omitted whenever it is clear from the context that quantities from the current period are referred to. There are constant marginal production costs c_i for each firm $i = 1, 2$ and no fixed costs. Therefore, firm i's total production costs C_i take the form

$$C_i(q_i) = q_i c_i \tag{2}$$

---

(1) Each firm i is assigned a base quantity b_i.

(2) Each period, the firm randomly picks its q_i from the uniform distribution $b_i \pm \delta_i$, where δ_i is a firm-specific parameter. This is the quantity actually "produced" by the firm in the current period. The firm "probes" or "explores" the space around b_i for advantageous production quantities.

(3) "The market" collects the quantities from all the firms, calculates Q in the current period, and determines the current price P according to the linear demand function. Markets are assumed to clear in all periods.

(4) Firms calculate their profit $\pi_i = q_i(P - c_i)$.

(5) Each firm maintains two vectors: π_i^a for profits generated when the current quantity is above the base quantity (i.e., $q_i \geq b_i$) and π_i^b for profits generated when the current quantity is below the base quantity (i.e., $q_i < b_i$).

(6) Each firm repeats steps 2 -5 for an entire *epoch*, which typically consists of 30 periods. By the end of an epoch:

   (i) The firm calculates the averages of the profits generated when the quantity was above, respectively below the base quantity:

$$avg_i^a = \frac{\sum \pi_i^a}{\dim \pi_i^a}, avg_i^b = \frac{\sum \pi_i^b}{\dim \pi_i^b}$$

   (ii) If $avg_i^a \geq avg_i^b$, the firms makes an upward adjustment of its base quantity b_i for the next epoch, i.e., $b_i = b_i + \epsilon_i$, where ϵ_i is a firm-specific parameter.

   (iii) Otherwise, $b_i = \max(0, b_i - \epsilon_i)$

   (iv) Both vectors π_i^a and π_i^b are reset.

   (v) The firm continues as normal at step 2.

Figure 1: "Probe and Adjust" according to Kimbrough and Murphy (2009, p. 56).

The firms decision procedure termed "Probe and Adjust" is described in Figure 1 (cf. Kimbrough and Murphy 2009, Figure 1, p. 56). As Kimbrough and Murphy point out, "Probe and Adjust" is in line with other algorithms where agents find the "direction of improvement" (Winston, 2004). It is worth noting that here the firms quantities are continuous, while, as we will see later, Huck and Oechssler (2000) work with a "finite grid".

As can be inferred from Figure 1, a firm has several "inherent" properties, which might be identical or different for other firms. For clarity and further use, these properties are summarized again in Figure 2, together with default values as used by (Kimbrough and Murphy, 2009, cf. Table 2, p. 60). Unless otherwise stated, these default values are also used in the analysis in this paper.

Base quantity b_i: The initial quantity to start the process. As it is a priori unclear whether there are any initial value effects in this algorithm, firms might have different initial quantities. **Default value for both firms**: 40.

δ_i**:** Determines the range within which the firm randomly draws its production quantity in each period. As the production quantity is $b_i \pm \delta_i$, the range is $2\delta_i$. **Default value for both firms**: 3.

ϵ_i**:** Determines by which amount the base quantity b_i should be adjusted by the end of an epoch. **Default value for both firms**: 0.7.

Epoch length: For how many periods the firm continues exploring possibilities around the current b_i until it adjusts it. **Default value for both firms**: 30.

Unit costs c_i: The constant marginal cost of production of one unit of the final good. **Default value for both firms**: 0.

Figure 2: Properties of the firms and their default values.

By default, the intercept for Equation (1) is $\alpha = 400$, and the slope is $\beta = 2$. At unit costs $c_i = 0$ for both firms in the duopoly, the Cournot quantity is therefore

$$Q^C = \frac{2\alpha - c_1 - c_2}{3\beta} = \frac{2\alpha}{3\beta} = 133.33 \tag{3}$$

If the market were fully competitive and marginal costs were equal among firms, i.e., $c = c_1 = c_2 = 0$, the Walrasian outcome would be

$$Q^W = \frac{\alpha - c}{\beta} = \frac{\alpha}{\beta} = 200 \tag{4}$$

Finally, a monopolist operating in this market with constant marginal costs $c = c_1 = 0$ would choose the production quantity

$$Q^M = \frac{\alpha - c}{2\beta} = \frac{\alpha}{2\beta} = 100 \tag{5}$$

The reference values $Q^C = 133.33, Q^W = 200$ and $Q^M = 100$ are used as benchmarks in the following sections.

3.2 Huck and Oechssler (2000)

Huck and Oechssler discuss a Cournot oligopoly in purely analytical framework. Even though their results are more general, for reasons given above, I here only describe the case with two firms, further on being referred to by an index $i = 1, 2$. Firms have no information about the decisions of the competitors, but they observe which price they achieve in each period, and are able to calculate their profits consequently. They have a two-period memory, which allows them to store their previous period's quantity and profit. Firms are assumed to optimize absolute profits by following a learning process called "trial and error". In each period t, each firm i chooses its production quantity $q_{i,t}$ from the finite grid $\Gamma := \{0, \delta, 2\delta, \dots, v\delta\}^{2,3}$. Note that this is a fundamental difference to Kimbrough and Murphy, where firms pick their quantities from continuous space. However, also in Huck and Oechssler, the choice of $\delta > 0$ is arbitrary, and $v \in \mathbb{N}$ should be "large enough" (Huck and Oechssler, 2000, p. 2). The inverse demand function $P_t(Q_t)$ is assumed to satisfy $P'_t < 0$ and $P'_t + 2P''_t Q_t < 0$, where $Q_t := q_{1,t} + q_{2,t}$ as in Section 3.1. This assumption is satisfied by the linear inverse demand function described in Equation (1). Production costs for each firm are assumed to be increasing and weakly convex, which again is satisfied by working with constant marginal costs $c_i > 0$ for each firm i as in Equation (2). In order to obtain their results analytically, they require existence of some Q_t such that $P_t(Q_t) = C'_i(0)$. By combining Equation (1) and Equation (2), this is clearly satisfied for $q_i = 0$ and $q_j, j \neq i$, large enough such that $P_t(Q) = 0$. Finally, "to avoid a monopolized market" (Huck and Oechssler, 2000, p. 2), it is assumed that the price at which each firm would sell if it were a monopolist is larger than the minimal marginal cost of all other firms. Since in Equation (2) $C'_i(0) = 0$, also this assumption holds.

For choosing production quantities $q_{i,t}$, firms follow a process called "trial and error", which is summarized in a simple formula:

$$d_{i,t} = \text{sign}(q_{i,t-1} - q_{i,t-2}) \times \text{sign}(\pi_{i,t-1} - \pi_{i,t-2})$$

$$s_{i,t} = \begin{cases} d_{i,t} & \text{if } d_{i,t} \neq 0 \\ \frac{1}{3}[-1] + \frac{1}{3}[0] + \frac{1}{3}[-1] & \text{if } d_{i,t} = 0 \end{cases}$$

$$q_{i,t} = q_{i,t-1} + \delta s_{i,t} \tag{6}$$

where as before $\pi_{i,t} = q_{i,t}(P_t - c_i)$ and $\text{sign}(x)$ is defined as

$$\text{sign}(x) = \begin{cases} 1 & \text{if } x > 0 \\ 0 & \text{if } x = 0 \\ -1 & \text{if } x < 0 \end{cases}$$

In period t, firm i starts its consideration with the quantity chosen in the previous period $q_{i,t-1}$. Depending on the direction of the quantity change and its respective profit in the previous period, the quantity in the current period is one step δ above or below the previous periods quantity. If

---

[2] As in Section 3.1, the subscript t is omitted if it is clear from the context which period this is about.

[3] The δ used here could be any positive value. For simplicity and notational convenience, I further on assume that $\delta = \delta_1 = \delta_2$ from the previous section.

in the previous period (i) the quantity was increased and the profit consequently increased as well (both signs being positive yields a positive factor for the δ) or (ii) the quantity was decreased and the profit decreased as well (both signs being negative yields a positive factor for the δ as well), the quantity is increased in the current period. Otherwise, the quantity is decreased in the current period. This process represents the idea that if last period's move was profitable, the direction of the move is kept; otherwise, it is reverted.

If any of the two sign expressions is 0, e.g., because both firms produce such that $P_t = c_i$, resulting in zero profits, then the firms would get stuck according to Equation (6). Hence, there is the extra rule that if $d_{i,t} = \text{sign}(q_{i,t-1} - q_{i,t-2}) \times \text{sign}(\pi_{i,t-1} - \pi_{i,t-2}) = 0$, the firm chooses a direction of movement with equal probability $\frac{1}{3}$ [4].

According to Huck and Oechssler, the process can be started with arbitrary values for $q_{i,t-1}$ and an arbitrary sign for the direction of movement of the previous period, or, equivalently, simply with two arbitrary $q_{i,t-1}, q_{i,t-2}$ and according profits.

Plugging in marginal production costs $c_1 = c_2 = 10$ into Equations (3) to (5) yields a Cournot quantity $Q^C = 130$, a Walrasian outcome $Q^W = 195$ and a monopoly quantity $Q^M = 97.5$.

4 Replication

Based on the models described in the previous section, I set up my own model to match the reference models as close as possible. This section describes to what extent I am able to replicate the results found in the original papers. For my model I use the simulation software Repast Simphony[5].

Wilensky and Rand (2007) propose to start with specifying a replication standard. Axtell et al. (1996) distinguish between aiming for (i) numerical identity, (ii) distributional equivalence and (iii) relational alignment. Numerical identity means that *exactly* the same values are produced by the replicated model as by the original model. However, even running exactly the same model on the same computer might result in slightly different results due to different floating point representation in memory (Belding, 2000). As numerical identity is also not mandatory for the analysis at hand, I do not aim for it. On the other extreme, relational alignment is satisfied if two models predict an output value in the same direction if an input parameter is changed in the same direction. This might not be sufficient for my needs, because I do not only want to know whether, for example, increasing the production costs for firm i results in lower production quantities of the firm, but rather whether, *on aggregate*, firms find the Cournot quantity. Finally, distributional equivalence means that both the original and the replicated model produce a distribution of one or more output values which cannot statistically be distinguished from each other. In terms of precision , this is between "numerical identity" and "relational alignment". Aiming for distributional equivalence should be manageable, yet provide all the key insights required. Thus, this is my replication standard for the model of Kimbrough and Murphy (2009). Since the model of Huck and Oechssler (2000) is analytical and gives an exact result, another statistical test is used to check whether the distribution of the results provided by my model is statistically distinguishable from the target

---

[4]Huck and Oechssler in principle allow firm randomly experimenting with some probability ϵ by choosing any direction of change. However, the key theorem holds only for the limit case where $\epsilon \to 0$. Therefore, experimenting is left aside for the moment.

[5]http://repast.sourceforge.net/, in the latest stable version 2.1.

value. For both models, the target measure is the total production quantity $Q := q_1 + q_2$. Some of the results obtained by Kimbrough and Murphy (2009) go far beyond the key question addressed here, and therefore is not be replicated.

Unless otherwise stated, I work with a linear inverse demand function as in Equation (1), a linear cost function as in Equation (2), and the set of default values for firms behavior as in Figure 2.

4.1 Kimbrough and Murphy (2009)

Kimbrough and Murphy repeat the simulation with the above described default values 100 times and describe the results in their paper. However, in order to perform nonparametric statistical tests between their and my results, I had to run their simulation. It is not possible to conduct a reasonable statistical test simply by comparing the means and / or the minimum and maximum values. A full set of data is required. As Kimbrough and Murphy generously provide their NetLogo model online, this was an easy task. Thus, the exact values that are further on referred to as values from Kimbrough and Murphy are **not** exactly the ones they used in their paper, but the ones that result from running their simulation on my own.

In one simulation run, Kimbrough and Murphy finish the simulation after 6600 periods of "probe and adjust", which, as they claim, is "long after a stable settling has occurred" (Kimbrough and Murphy, 2009, p. 58). They then calculate the average of the total quantity Q over the last 1000 periods (in their terminology, this quantity is called "$runningAverageBid$"). Table 1 contains a summary of the 100 "average total quantities" obtained from their model, compared with the respective values as obtained from my model. Keeping in mind that for the default parameter settings with $c_1 = c_2 = 0$ the total Cournot quantity is $Q^C = 133.33$, the data provided is a reasonable basis for their key finding that "Probe and Adjust with Own Returns leads the agents to the Cournot solution"[6] (Kimbrough and Murphy, 2009, p. 60)[7].

	Kimbrough Avg. Q	My Avg. Q
Min.	129.9	129.6
1st Qu.	132.7	132.3
Median	133.6	133.4
Mean	133.6	133.3
3rd Qu.	134.7	134.4
Max.	137.2	137.9

Table 1: Average quantities obtained from Kimbrough and Murphy (2009) and my model.

The output of a representative simulation run is shown in Figure 3. The data suggests that there is no statistical significant difference between the results obtained by Kimbrough and Murphy and my model. The nonparametric Mann-Whitney U test (also referred to as Wilcoxon rank sum test) was used to check for the null hypothesis that the distributions of the two results do not differ. The p-value of the Mann-Whitney U test is 0.294; thus, there is no evidence to reject the null hypothesis that the results do not differ. Therefore, the model replication can be considered

---

[6]As mentioned above, "Own Returns" means absolute profit maximization.

[7]Kimbrough and Murphy base their claim on the result of a binomial test. This test only captures how many outcomes where above or below a certain value, but it does not take into account how far the outcomes were from the target value. Therefore in the following, the binomial test is not used, but the more appropriate t-test instead.
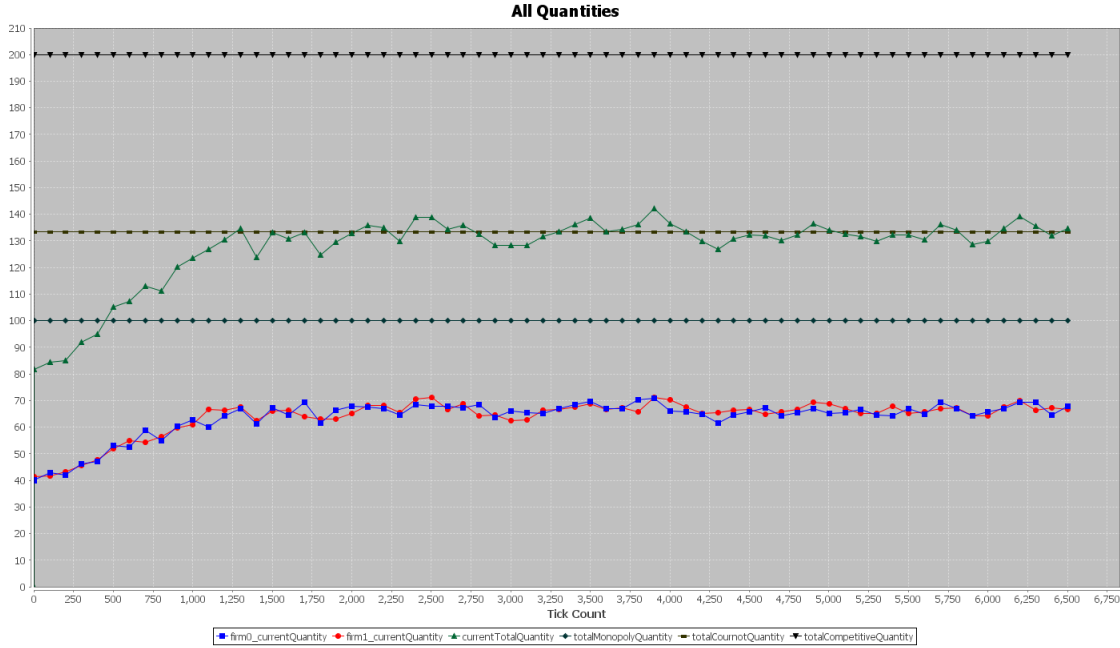
Figure 3: Firms behaving according to "probe-and-adjust" with default parameter values.

successful; the results obtained from my model are distributionally identical to the results of Kimbrough and Murphy.

4.2 Huck and Oechssler (2000)

For a model setup as described above with the default parameters from Figure 2 with constant marginal costs $c_1 = c_2 = 10$, Huck and Oechssler predict that the firms settle at the "joint profit maximizing outcome", i.e., $Q := q_1 + q_2 = Q^M = 97.5$.

Setting up an agent-based model with the firms behaving according to "trial and error" as described in Section 3.2 is straightforward. In the first two periods, firms choose any random quantity between 0 and their respective base quantity, here being set to 400 for both firms. This setup satisfies the requirement of the model for an "arbitrary initial quantity". The step size of both firms δ is set to 0.5, in order to allow "finding" exactly the expected outcome of $Q^M = 97.5$. As first experiments showed, firms tend to find stable behavior way faster than with "probe and adjust"; on average, after 1200 - 1500 periods. Thus, each simulation run does not last for 6600 periods, but only 3000 periods, still way after stable behavior has established. Again, the total quantity averaged over the last 1000 runs of the simulation is taken and reported. I repeated the simulation run 200 times. For each of those 200 runs, the average total quantity was exactly 97.5, with no exceptions. Thus, there is not even a need for a statistical test; clearly, firms in an agent-based simulation behaving according to "trial-and-error" find the "joint profit maximizing outcome", as predicted by Huck and Oechssler. Thus, also this model replication can be considered successful.

The exceptional rule that, if in Equation (6) any of the two sign-expressions is equal to 0, the firms "experiment" with some random probability is of utmost importance. Without that rule,

firms get stuck at some quantities, having no possibility to deviate from there again.

It is also worth noting that firms are not right on Q^M all the time; instead, they keep oscillating around it. This is inherent to the nature of "trial and error".

5 From Kimbrough and Murphy towards Huck and Oechssler

Having successfully replicated both the model of Kimbrough and Murphy and of Huck and Oechssler, including the contradicting results of finding the Cournot quantity and the monopoly quantity, respectively, it is now time to attempt to align the two models with each other, in order to find out which ingredient of the models causes the striking difference. Both models make their firms start with a base quantity, from where on they attempt to improve. Once an improving quantity is found, future possibilities are explored from there on. So what exactly causes the differing outcomes? In the following paragraphs, stepwise changes are introduced to the replicated model of Kimbrough and Murphy.

For the replication of Kimbrough and Murphy as in Section 4.1, marginal production costs of 0 were used, i.e., $c_1 = c_2 = 0$. However, Huck and Oechssler work with increasing cost functions, so for my replication in Section 4.2 I used $c_1 = c_2 = 10$, resulting in $Q^C = 130$ and $Q^M = 97.5$. Keeping the default parameter values from Figure 2, but using $c_1 = c_2 = 10$ and repeating a simulation with 6600 periods 200 times results in a distribution of average total quantities over the last 1000 periods of the run as shown in Table 2. The t-test for the null hypothesis that the mean of the distribution is $Q^C = 130$ cannot be rejected ($p = 0.413$). Thus, one can conclude that "probe and adjust" is, as expected, robust to introducing positive marginal production costs.

	Avg. Q
Min.	125.8
1st Qu.	128.8
Median	130.2
Mean	130.1
3rd Qu.	131.3
Max.	134.1

Table 2: Average quantities obtained from my "probe and adjust" model with $c_1 = c_2 = 10$ and $Q^C = 130$.

An obvious difference between the setup of the two models is that while Huck and Oechssler let their firms choose their quantities from a finite grid, Kimbrough and Murphy use continuous space. Thus, a modified version of my "probe and adjust" model makes firms not pick any more from the uniform distribution $b_i \pm \delta_i$ as in Step 2 of Figure 1, but from $\{b_i - \delta_i, b_i, b_i + \delta_i\}$, each with probability $\frac{1}{3}$. However, the average total quantity remains effectively unchanged around the Cournot quantity of $Q^C = 130$. Therefore, this difference in the model setup cannot explain the differing outcomes.

Up to now, the simulation is run with the default values $\delta_i = 3, \epsilon_i = 0.7, i = 1, 2$, representing the range from which to pick quantities in each round (which is $2\delta_i$), and the quantity by which the base quantity b_i is adjusted by the end of an epoch, respectively. Using these parameters (and, for that matter, any combination of δ_i and ϵ_i with $\delta_i > \epsilon_i$), it is possible that, even though firm i *increased* its base quantity in the last round by 0.7, the quantity in the next period is effectively

below the previous quantity, even though the increase was apparently profitable (otherwise it would not have been chosen in the first place). This is at odds with the setup of Huck and Oechssler, where the quantity chosen in a period is always strictly *above* the quantity chosen in the previous period, if this increment was profitable. Thus, I re-run the "probe and adjust" simulation with $\delta_i = 0.5, \epsilon_i = 1, i = 1, 2$, closer resembling the setup of "trial and error" from Huck and Oechssler. Nevertheless, the average total quantity produced by the model remains around the Cournot quantity $Q^C = 130$. This verifies the results obtained by Kimbrough and Murphy, who find robustness of their model to such parameter changes.

A central element of "trial and error" is using the direction of change of previous quantities and profits. In a sense, there is a similar element in "probe and adjust" - if producing quantities above the base quantity is more profitable then producing below the base quantity, the firm is *more likely* to continue doing so in the future, but it is not *guaranteed* to do so; it simply rather explores future options from there on. In that sense, "probe and adjust" firms do not stick strictly to any direction of movement. A first step to make "probe and adjust" firms less fragile is by reducing the epoch length from the default value of 30 to 2. Introducing this change (on top of the other modifications already described), however, does not contribute to less fragility; on the contrary, the fluctuations get even bigger, and it is questionable whether total quantities still converge and if so, to which level. A representative output after 66000 periods - 10 times longer than the duration of the previous runs - is shown in Figure 4[8].
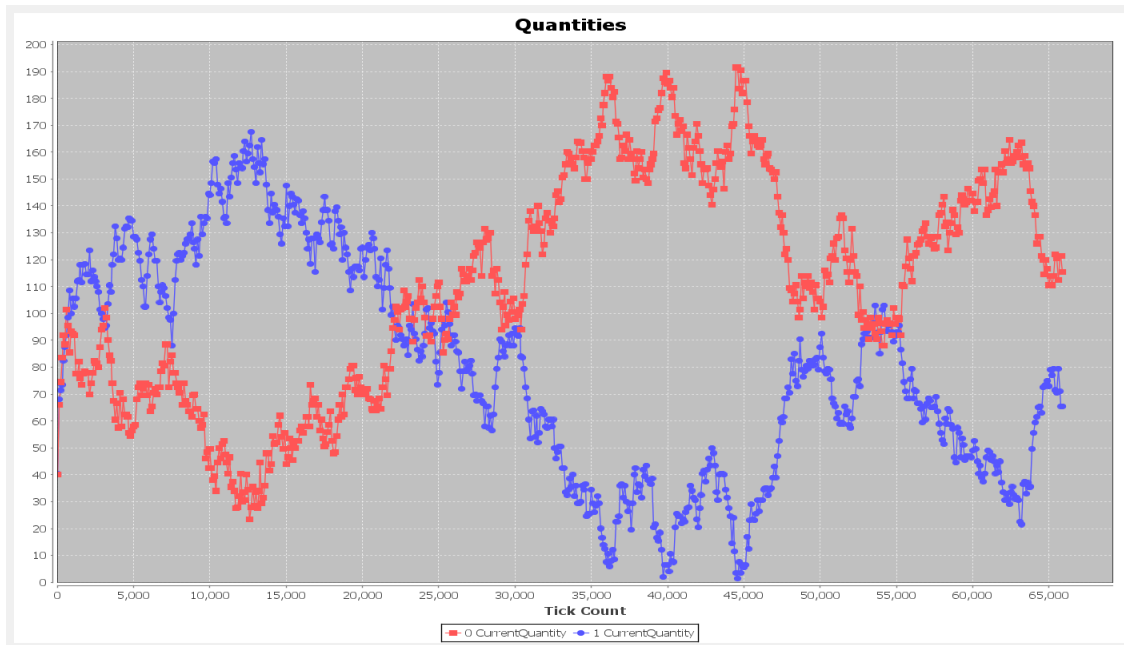

Figure 4: Firms behaving according to the modified "probe-and-adjust", epoch length = 2.

Now, by employing an epoch length of 2 in the current setup, there is a possibility that firms

---

[8]Kimbrough and Murphy note that while their original "probe and adjust" algorithm is robust, the standard deviation of the average total quantity increases with smaller epoch lengths (excluding an epoch length of 1 where "... agents engage in what looks like a random walk ... ", Kimbrough and Murphy 2009, p. 59). This robustness is apparently lost with making the firms decisions not on continuous space, but on a finite space, as described above. However, this is not the scope of the current analysis, and is therefore left aside.

make a move in the same direction twice, not allowing them to compare the results with the alternative direction. In a sense, they lack the ability to compare the results as "trial and error" firms do. Thus, keeping an epoch length of 2, firms "probe and adjust" behavior is modified again as follows:

- In the first round of an epoch, firms choose from the uniformly distributed continuous space $b_i \pm \delta_i$ again, as apparently using finite space instead was not of much help.

- In the second round, firms play the *opposite* direction from the first round. If in the first round the chose some value $q_{i,t-1} > b_i$, they now choose $q_{i,t} = \max(0, 2b_i - q_{i,t-1})$, and vice versa. This causes them to play exactly once *above* its base quantity b_i, and once below. The only exception is that if they played b_i in the first round, they repeat that decision.

- The adjustment rule by the end of the epoch, i.e., after the second round, remains unchanged, as described in Figure 1.

Due to the randomness in the first round of each epoch, still getting stuck, as would be possible for "trial and error" firms without the exceptional rule, should be avoided. This behavior should enable "probe and adjust" firms "memorize" two decisions in opposing directions and to subsequently compare the profits generated respectively, as is done by "trial and error" firms. Table 3 compares summary statistics from 200 repeated simulation runs of my unmodified "probe-and-adjust" model with runs from the newly modified version. The standard deviation of the average total quantity has significantly increased, but the average total quantity is still around the Cournot quantity $Q^C = 130$. However, the results are still different from the monopoly quantity $Q^M = 97.5$ found by Huck and Oechssler.

	p-a-a avg. Q	p-a-a avg. Q sd	p-a-a-m avg. Q	p-a-a-m avg. Q sd
Min.	125.8	2.59	122.7	5.37
1st Qu.	128.8	2.91	128.3	6.90
Median	130.2	3.11	130.7	7.95
Mean	130.1	3.12	130.7	8.14
3rd Qu.	131.3	3.34	132.7	9.02
Max.	134.1	4.47	139.4	12.79

Table 3: Average total quantities obtained from "probe and adjust" (p-a-a) and "modified probe and adjust" (p-a-a-m) and their standard deviations (sd).

It seems that the firms still lack some kind of continuity, both on their side and on the side of their competitor. "... [firms can do better than the Cournot outcome because] further simultaneous downward adjustments still increase the profits of both firms. And this continues until they reach the joint-profit maximum" (Huck et al., 2004a, p. 206). Even the modified "probe and adjust" firms might try to explore the space *above* its base quantity in the next period; and if not the firm itself, still its competitor may do so. It is much more unlikely that both of them decide to decrease their quantities and thus find a mutually beneficial result, and if they happen to produce in opposing directions, they both find it (misleadingly) beneficial to increase their quantities - until they end up in the Cournot outcome again. Therefore, I changed the "probe and adjust" algorithm once more. Instead of randomly picking some quantity in the first and in the second period of an epoch, firms now follow a much more straightforward rule:

- In the first period of an epoch, $q_i = b_i + \delta_i, i = 1, 2$. This means that both firms attempt to increase their production quantities above the base quantities.

- In the second round, $q_i = \max(0, b_i - \delta_i), i = 1, 2$, i.e., firms decrease their production quantities below the base quantities.

- The adjustment rule by the end of the epoch, i.e., after the second round, remains unchanged, as described in Figure 1.

For initial quantities $b_i = 40, \delta_i = \epsilon_i = 0.5, i = 1, 2$, the firms choose the joint-profit maximizing outcome after less than 200 periods. This model is not at all random any more, but fully deterministic. Thus, in order to exclude the possibility of dependence on the initial quantities, both firm's initial value is set to a random value between 0 and 400, which matches the choice of the initial values in Section 3.2. For 200 repeated simulation runs, each lasting for 3000 periods, the total production quantity, averaged over the last 1000 periods, equals the monopoly quantity $Q^M = 97.5$ in each run. This shows that finally this strongly modified version of "probe and adjust" is able to reproduce the results of "trial and error" firms. It seems that the key difference is randomness in the behavior of "probe and adjust" firms.

In order to prove that conjecture, I also modified the "probe and adjust" algorithm to make firms randomly deviate from their default behavior with a propensity to experiment of $\gamma_i \in [0, 1], i = 1, 2$[9]. If firms choose to experiment, they pick each from $\{q_{i,t-1} - \delta_i, q_{i,t-1}, q_{i,t-1} + \delta_i\}$ with probability $\frac{1}{3}$, as in the standard exceptional rule of Huck and Oechssler (see Section 3.2). If $\gamma_1 = \gamma_2 = 0$, the behavior naturally remains unchanged. For the other extreme case with $\gamma_1 = \gamma_2 = 1$, there is no rigidity at all and firms make totally random decisions in each period, consequently finding no convergence towards any stable combination. Both firms benefit from the other firm sticking to its direction instead of experimenting.

Concluding, the missing link between the original version of "probe and adjust" by Kimbrough and Murphy and "trial and error" by Huck and Oechssler is rigidity. "Trial and error" firms are able to converge to the joint-profit maximizing quantity by not only sticking to their direction of movement, but also by profiting from the fact that their competitor also sticks to his direction of movement. (Modified) "Probe and adjust" firms benefit from the fact that both of them always increase their production quantity in the first period, and always decrease it in the second period. Both algorithms find a the joint-profit maximizing outcome without any explicit coordination device among the firms. In fact, firms are able to benefit from the rigid production behavior of their opponents, even without anticipating them or taking them into account. These findings show that the randomness involved in the original version of "probe and adjust" prevents further improvements beyond the Cournot quantities for both firms.

6 Sensitivity analysis

In the previous section, I showed that the propensity to experiment γ_i is a crucial parameter describing the behavior of "trial and error" firms. In this section, I explore for which combinations

---

[9]In the notation of Huck and Oechssler, this parameter is called ϵ. For clear distinction from the parameter ϵ already used for "probe and adjust" firms in Section 3.1, I refer to this parameter as γ instead.

of $\gamma_i, i = 1, 2$, the outcome of convergence to the total monopoly quantity is robust. I also show that, even if solutions still converge to the total monopoly quantity, stability is reached much later for firms with a higher propensity to experiment.

Firm i with a propensity to experiment γ_i chooses the production quantity in period t, $q_{i,t}$, according to the standard "trial and error" rule in Equation (6) with probability $1 - \gamma_i$. However, with probability γ_i, the firm decides to experiment instead, and chooses a random direction of movement. Therefore "trial and error" with experiments can be described as follows:

$$d_{i,t} = \text{sign}(q_{i,t-1} - q_{i,t-2}) \times \text{sign}(\pi_{i,t-1} - \pi_{i,t-2})$$

$$s_{i,t} = \begin{cases} (1 - \gamma_i)[d_{i,t}] + \gamma_i(\frac{1}{3}[-1] + \frac{1}{3}[0] + \frac{1}{3}[+1]) & \text{if } d_{i,t} \neq 0 \\ \frac{1}{3}[-1] + \frac{1}{3}[0] + \frac{1}{3}[+1] & \text{if } d_{i,t} = 0 \end{cases}$$

$$q_{i,t} = q_{i,t-1} + \delta_i s_{i,t} \tag{7}$$

$d_{i,t}$ is the direction of movement from the previous period, taking into account whether this move was beneficial or not. The case with $d_{i,t} = 0$ reflects the exceptional rule described above, meaning that firms choose a random direction of movement in order to prevent getting stuck.

The parameter space

$$(\gamma_1, \gamma_2) = \{(0.0, 0.0), (0.05, 0.05), (0.1, 0.1), (0.2, 0.2), (0.3, 0.3), (0.4, 0.4),$$
$$(0.5, 0.5), (0.6, 0.6), (0.7, 0.7), (0.8, 0.8), (0.9, 0.9), (1.0, 1.0)\}$$
$$\cup \{0\} \times \{0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, 1.0\}$$

is analyzed, i.e., combinations where none of the firms ever experiments, both of them always experiment, some symmetric combinations in between as well as some cases where firm 1 never experiments, but firm 2 experiments sometimes. For each combination, the simulation is repeated 100 times, with the other parameter values as before: $c_1 = c_2 = 10$, resulting in $Q^C = 130$ and $Q^M = 97.5$. In order to stay in line with the simulation conducted by Kimbrough and Murphy, each simulation lasts for 6600 periods, which in their setup is "long after a stable settling has occurred" (Kimbrough and Murphy, 2009, p. 58). The following numbers are used for comparison between the results: As before, the (i) average of the total production quantity ("average total quantity") over the last 1000 periods of the simulation are calculated by the end of the simulation, as well as (ii) the standard deviation of the total production quantity over the last 1000 periods ("total quantity standard deviation"). In order to measure the adjustment speed towards the total monopoly quantity Q^M, (iii) the period in which the total monopoly quantity Q^M is reached for the first time is recorded. In the following, this is referred to as "convergence period". If Q^M is never reached during an entire simulation run, this simulation run is referred to as non-converging. Thus, for each combination of (γ_1, γ_2), one can calculate (iv) which percentage of the simulation runs was converging.

Figure 5 and Figure 6 give an overview about the simulation results. The numbers represent the average of the respective values, calculated over the 100 repeated simulation runs for this
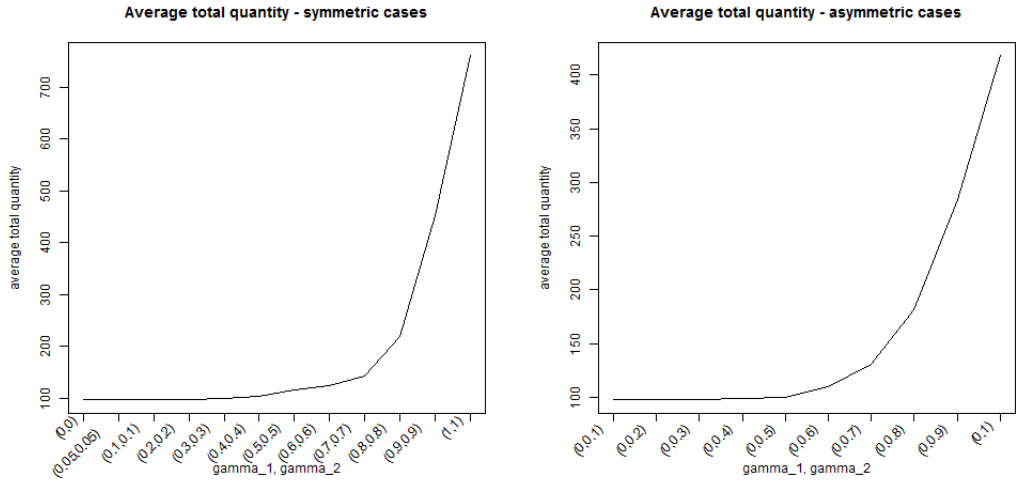
(a) Average of average total quantity - symmetric cases

(b) Average of average total quantity - asymmetric cases
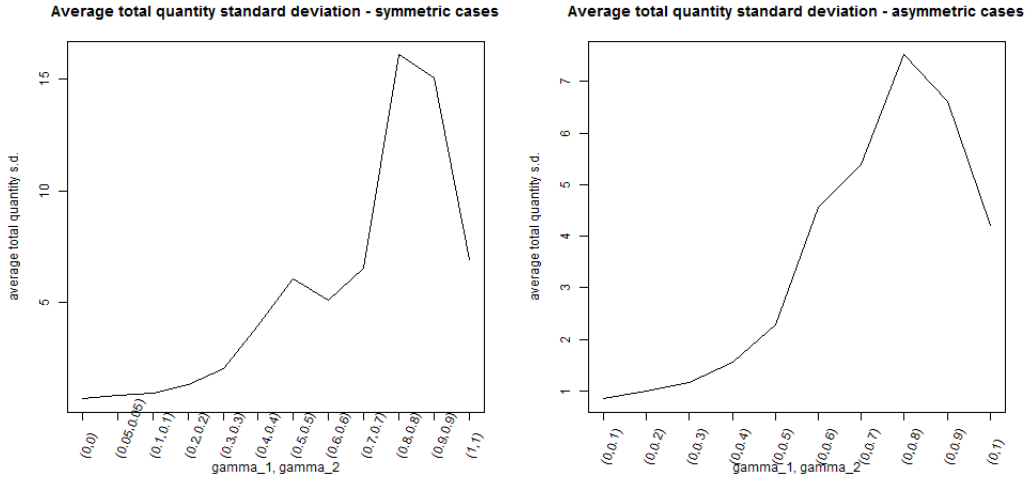
(c) Average of standard deviation of total quantity - symmetric cases

(d) Average of standard deviation of total quantity - asymmetric cases

Figure 5: Total production quantities for "trial and error" with experiments and different (γ_1, γ_2)

particular combination of (γ_1, γ_2)[10]. On all graphs, the x-axis shows the combinations of (γ_1, γ_2), and the y-axis the respective quantities of interest. First, different combinations with simultaneous increment of the propensity to experiment are explored, from $(0,0)$ to $(1,1)$. Combinations further to the right represent higher joint propensity to experiment, until the combination $(1,1)$. Next combinations are explored where only *one* firm experiments,, i.e., $\gamma_1 = 0$ and γ_2 increases stepwise. So again, the propensity to experiment of firm 2 increases from the left to right from $(0, 0.1)$ to $(0, 1)$. Figure 5(a) and Figure 5(b) show the average of "average total quantities". Figure 5(c) and Figure 5(d) show the average of the standard deviation of the total production quantity (*not the standard deviation of the average total production quantity*). Figure 6(a) and Figure 6(b) show which percentage of the 100 simulation runs found any convergence at all, i.e., how many of the
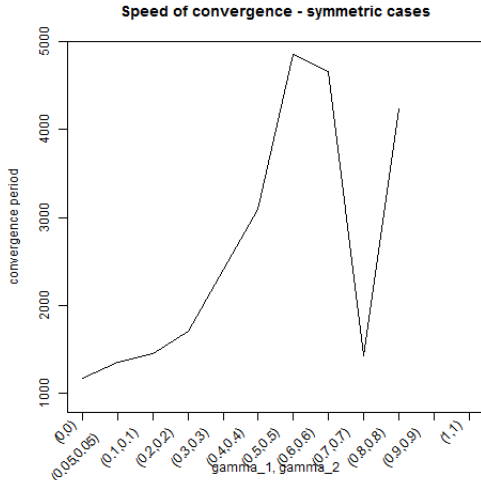
_____
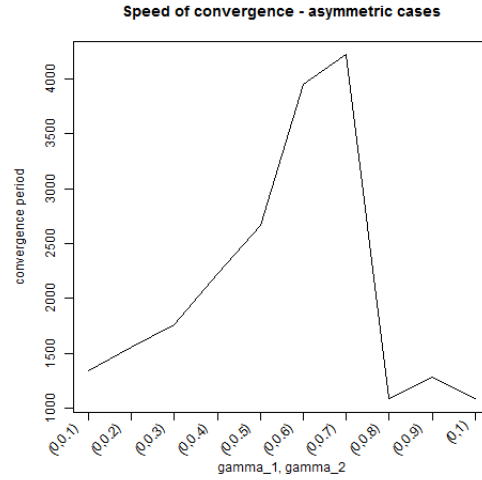[10]The exact numbers are provided in Table 5 in the appendix.
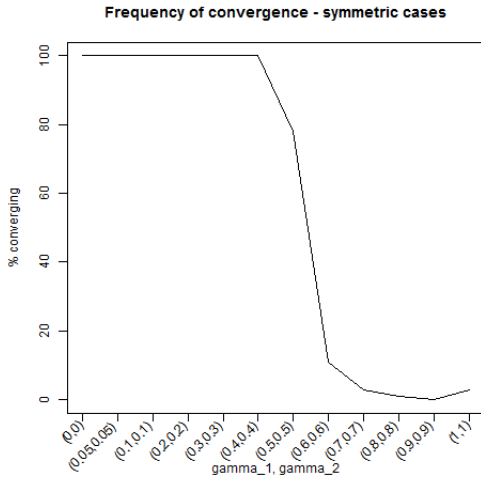
(a) Percentage of converging simulation runs - symmetric cases
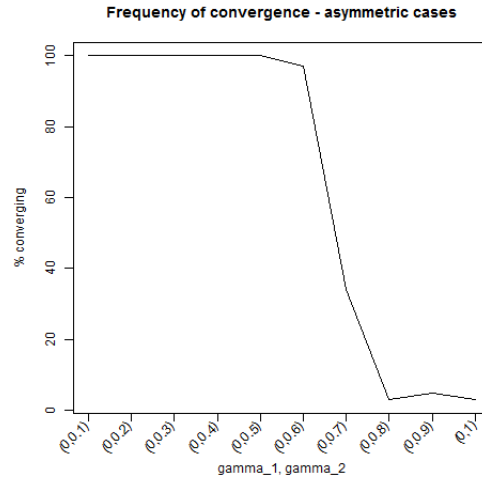
(b) Percentage of converging simulation runs - asymmetric cases

(c) Average first convergent period - symmetric cases

(d) Average first convergent period - asymmetric cases

Figure 6: Convergence for "trial and error" with experiments and different (γ_1, γ_2)

simulation runs ever reached a total production quantity of $Q^M \pm 5\%$. Figure 6(c) and Figure 6(d) show the average first period with convergence to Q^M, for those simulation runs which reached it all, i.e., which were considered "convergent" in the first place. A lower number thus represents faster convergence.

The combination $\gamma_1 = \gamma_2 = 0$ is the benchmark result, where firms, as in the initial "trial and error", do not experiment at all. In 100% of the simulation runs, they reach the total monopoly quantity $Q^M = 97.5$, on average within 1175 periods. On the other extreme, the combination $\gamma_1 = \gamma_2 = 1.0$ represents totally random behavior of the firms. Even though they come close to Q^M in 3 out of 100 runs, the average total production quantity at 760 shows that there is no more any meaningful behavior. All other combinations settle somewhere between the two extremes.

It is apparent that a higher propensity to experiment for only one firm has way more drastic

implications than a moderate propensity to experiment for both firms. For instance, for $\gamma_1 = \gamma_2 = 0.3$, the average total production quantity is still only 2% above Q^M, whereas for $\gamma_1 = 0, \gamma_2 = 0.6$, the average total production quantity is already 13% above Q^M. For both combinations, firms take much longer to converge to Q^M, if at all, than non-experimenting "trial and error" firms. For the first combination, the value is already twice the benchmark value of 1175 periods, and for the second, it is even three times the benchmark value.

The data suggests that "trial and error" is relatively robust to moderate changes in the propensity to experiment with respect to average total production quantity, as shown in Figure 5(a) and Figure 5(b). Even for propensities to experiment as high as $(0.4, 0.4)$, i.e., firms choose to experiment 4 out of 10 times instead of sticking to what they know to be beneficial for them already, the average total production quantity is only 6% above Q^M. Only for values above $(0.5, 0.5)$ the average total quantity is closer to the Cournot quantity $Q^C = 130$ than to Q^M. As mentioned above, high propensities to experiment of only one firm causes the outcome to drift away from Q^M much faster.

While the average total quantity is relatively robust to higher propensities to experiment, fluctuations naturally become much higher the more firms tend to experiment (Figure 5(c), Figure 5(d)). Even for $(0.2, 0.2)$, the average of the standard deviation of the total quantity is already twice as high as for $(0, 0)$, and for $(0.5, 0.5)$ even nine times higher than for $(0, 0)$. Again, versatility is higher if only one firm shows more random behavior than if both firms have moderate random behavior.

This is also reflected in the frequency of convergence (Figure 6(a) and Figure 6(b)), and if there is convergence at all, in the speed of convergence. While firms until combinations of propensities to experiment of $(0.4, 0.4)$ always find Q^M sooner or later, the likelihood drops quickly from there on: 78% for $(0.5, 0.5)$, 11% for $(0.6, 0.6)$, and 3% for $(0.7, 0.7)$[11]. If there is convergence at all, firms tend to find it much later if they experiment (Figure 6(c) and Figure 6(d)). Firms who experiment with $\gamma_1 = \gamma_2 = 0.5$ require, on average, 4.5 times longer to find Q^M for the first time than firms who do not experiment at all.

7 Relative profit maximization

7.1 Vega-Redondo (1997)

In his seminal paper, Vega-Redondo (1997) models a Cournot oligopoly in an explicitly dynamic way. Contrary to the traditional model setup, firms do not maximize absolute payoffs, but relative payoffs instead, by imitating the previous period's most successful firm. In order to imitate, firms require perfect knowledge about their competitors.

"Probe and adjust" firms, on the other hand, have extremely limited information about the environment in which they operate. In the model of Vega-Redondo, firms with perfect information converge to the fully competitive outcome Q^W. Does the same hold for "probe and adjust" firms, which have only imperfect information, when they maximize relative profits? Or do the firms in the model of Vega-Redondo exhibit some kind of "information curse", i.e., the additional information

---

[11]The fact that this value is 0% for $(0.9, 0.9)$, but somewhat surprisingly higher (3%) for $(1, 1)$ can be attributed to randomness.

they acquire induces a downward spiral in which all of them are worse off then they were without that information? In order to analyze that question, I first describe the setup of Vega-Redondo and then modify "probe and adjust" firms to maximize relative profits.

Relative profit maximization reflects the idea that survival is the primary consideration of firms, and if their competitors get stronger and stronger, they might be able to eventually exhibit market power. Alternatively, one might also "view [relative profit maximization] as responding to forces of learning and imitation" (Vega-Redondo, 1997, p. 382). However, this objective results in "spiteful behavior" of the firms, i.e., the willingness to take some actions which are harmful for themselves, as long as it harms their competitors even more.

Vega-Redondo models a n-firm oligopoly with a homogeneous product. Similar to Huck and Oechssler (2000), the firm's quantities are chosen from a finite grid, contrary to the continuous production space in the model of Kimbrough and Murphy (2009). The model is analyzed in discrete time.

Firms behavior is very simple and straightforward. In each period, firms are allowed to modify their previous production quantities with probability p. If the firm changes its production quantity, it either "imitates" the behavior (i.e., the production quantity) of last period's most successful firm, or with a small probability ϵ randomly chooses to experiment with any other quantity. Experimenting can be understood either as a firm's voluntary decision to experiment, "or they are replaced by some newcomer that chooses its output from *tabula rasa*" (Vega-Redondo, 1997, p. 379). In order to imitate a competitor, firms need to be able to perfectly observe their production quantities, which is a rather strict assumption. There is no description about the initial state of the system, suggesting that the results do not depend on the initial values and therefore are robust to arbitrary initial production quantities (which is again in line with Huck and Oechssler).

For any finite n, so in particular also for $n = 2$, total production quantities converge to the Walrasian outcome Q^W. Firms in the model of Vega-Redondo have full information and optimize *relative* instead of *maximize* profits. Thus, as prices and cost are the same across firms, the (relatively) most successful firm is the one with the highest production quantity. Therefore other firms are induced to increase their production quantities as well until no profits are made at all anymore at total production quantities Q^W. In the models of Huck and Oechssler ("Trial and error") and Kimbrough and Murphy ("Probe and adjust") firms have limited information and optimize maximize absolute profits, which allows them to do clearly better than in Q^W. In order to confirm that also firms trying to maximize relative profits that do not have full knowledge converge to the Walrasian outcome, I modify "Probe and adjust" accordingly.

7.2 "Probe and adjust" with relative profit maximization

As described in Section 3.1, "probe and adjust" firms are modeled to maximize absolute profit by default. Reflecting the idea of Vega-Redondo, I modify firms to optimize relative profit instead.

In order to calculate their relative profits, firms need to be aware of the overall market size, i.e., the total profits generated in a certain period. This is a slightly more restrictive information requirement than before, but still not too unrealistic. Firm i's relative profit in period t is then defined as

$$\pi_{i,t}^r \equiv \begin{cases} \frac{\pi_{i,t}}{\Pi_t} & \text{if } \pi_{i,t} \geq 0 \\ 0 & \text{if } \pi_{i,t} < 0 \end{cases} \tag{8}$$

where the total market profit is defined as $\Pi_t \equiv \sum_{i=1}^n \pi_{i,t}$. Afterwards, instead of comparing the average *total* profits (see Figure 1), firms memorize and compare average *relative* profits by the end of an epoch, leaving the remainder of the "probe and adjust" algorithm unchanged.

The second part of Equation (8) represents the idea that once a firm's absolute profit in a period is negative, it considers this as a relative profit of 0, disregarding the absolute profit of the other firms.

For the first run with the default parameter values (Figure 2), and in particular no production costs, i.e., $c_1 = c_2 = 0$, the result changes immediately drastically. Instead of converging to the total production quantity $Q^C = 133.3$, the total production quantity converges to $Q^W = 200$, as in the model of Vega-Redondo. However, as shown in Figure 7, contrary to "probe and adjust" with absolute profit maximization (Figure 3), the fluctuations are *not* symmetric around the target value Q^W. Once firms hit the total production quantity Q^W, they immediately find that their (relative) profits are higher below Q^W.
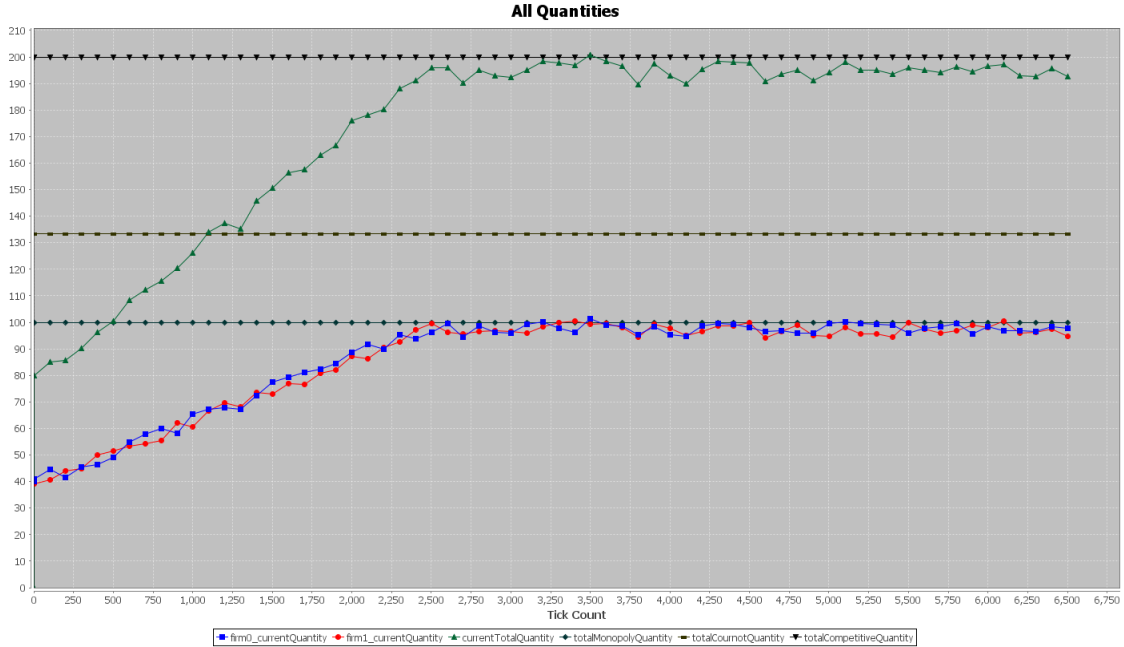
Figure 7: Firms behaving according to "probe-and-adjust", but maximizing relative profits, $c_1 = c_2 = 0$.

Summary statistics for 100 repeated simulation runs with $c_i = 0, i = 1, 2$ and $c_i = 10, i = 1, 2$ are shown in Table 4. The average total production quantities are close to Q^W, but clearly below it. As described above, this is due to the fact that once firms hit the ceiling of Q^W, they immediately decrease their production quantities. As there is very little space for fluctuations above Q^W, the mean of the process cannot be Q^W, but has to be slightly below it. Thus it is confirmed that also firms with limited information converge to a total production quantity (close to) Q^W, if they

	$c_1 = c_2 = 0$	$c_1 = c_2 = 10$
Min.	188.1	184.5
1st Qu.	193.3	188.5
Median	195.3	190.1
Mean	195.4	190.3
3rd Qu.	197.3	192.0
Max.	201.9	195.7
Q^W	200.0	197.5

Table 4: Average quantities from probe-and-adjust, maximizing relative profits.

optimize relative profits.

8 Conclusion

In this paper, I compare three models of oligopolistic firm behavior: "trial and error" by Huck and Oechssler, "probe and adjust" by Kimbrough and Murphy (having firms maximize absolute profits in both models), and a model with relative profit maximization through imitation by Vega-Redondo.

Even though the first two appear very similar at first sight, their results are at odds with each other. I find that the key component responsible for the difference in results is strict rigidity in the model of Huck and Oechssler versus exploration of the solution space, incorporating a random component, in the model of Kimbrough and Murphy. An argument frequently made in favor of randomization in agent behavior is accounting for unobservable decision criteria or erroneous decisions. According to my analysis, inclusion of a random component might change outcomes considerably. Moderate randomization may not have such drastic consequences, but definitely increases fluctuations. This might be important when discussing speed of convergence to stable behavior.

Similar to the model of Vega-Redondo, firms behaving according to "probe and adjust", having limited information, cannot do better than Q^W anymore, once they optimize relative profits instead of absolute profits. This is not very surprising, since once firms optimize relative profits, spiteful behavior, inherent to relative profit maximization, cannot be avoided. The model of Kimbrough and Murphy thus delivers the expected results if the target measure of the firms is modified.

Both for the inclusion of a random component as well as the appropriate target measure for firms, no clear recommendations can be given based on my analysis. The modeling decision should be made carefully, taking into account possible consequences, and with respect to economic intuition as well as on empirical grounds.

Appendix

"trial and error" with experiments

γ_1	γ_2	% convergence	convergence_period	average_total_quantity	average_sd
0.00	0.00	100	1174.89	97.50	0.71
0.05	0.05	100	1362.55	97.75	0.84
0.10	0.10	100	1449.53	97.98	0.98
0.20	0.20	100	1712.95	98.54	1.32
0.30	0.30	100	2412.14	99.60	2.05
0.40	0.40	100	3092.43	103.25	3.95
0.50	0.50	78	4844.51	115.45	6.07
0.60	0.60	11	4653.18	125.02	5.11
0.70	0.70	3	1425.67	143.65	6.54
0.80	0.80	1	4230.00	219.74	16.07
0.90	0.90	0		452.09	15.01
1.00	1.00	3	942.33	760.49	6.94
0.00	0.10	100	1344.69	97.76	0.85
0.00	0.20	100	1560.11	98.03	1.00
0.00	0.30	100	1759.70	98.31	1.16
0.00	0.40	100	2228.99	98.94	1.55
0.00	0.50	100	2666.70	100.40	2.29
0.00	0.60	97	3953.85	110.06	4.56
0.00	0.70	34	4217.62	130.45	5.39
0.00	0.80	3	1087.67	181.60	7.52
0.00	0.90	5	1280.00	284.38	6.60
0.00	1.00	3	1093.67	417.80	4.21

Table 5: Simulation results for 100 runs for each (γ_1, γ_2). All numbers represent averages over the 100 runs.

Bibliography

Arifovic, J. (1994). Genetic algorithm learning and the cobweb model. *Journal of Economic dynamics and Control*, 18(1):3–28.

Axtell, R., Axelrod, R., Epstein, J. M., and Cohen, M. D. (1996). Aligning simulation models: A case study and results. *Computational & Mathematical Organization Theory*, 1(2):123–141.

Barr, J. and Saraceno, F. (2005). Cournot competition, organization and learning. *Journal of Economic Dynamics and Control*, 29(1):277–295.

Baumol, W. J. and Quandt, R. E. (1964). Rules of thumb and optimally imperfect decisions. *The American economic review*, 54(2):23–46.

Belding, T. C. (2000). Numerical replication of computer simulations: Some pitfalls and how to avoid them. *Center for the Study of Complex Systems, University of Michigan*.

Bischi, G. I., Chiarella, C., and Kopel, M. (2002). On market games with misspecified demand functions: Long run outcomes and global dynamics. *University of Urbino; University of Technology, Sydney; University of Technology, Vienna*, September 2002.

Brandts, J. and Guillén, P. (2007). Collusion and fights in an experiment with price-setting firms and advance production. *The Journal of Industrial Economics*, 55(3):453–473.

Debreu, G. (1974). Excess demand functions. *Journal of mathematical economics*, 1(1):15–21.

Furth, D. (2009). Anything goes with heterogeneous, but not always with homogeneous oligopoly. *Journal of Economic Dynamics and Control*, 33(1):183–203.

Gravelle, H. and Rees, R. (2004). *Microeconomics*. Pearson Education, 3rd edition.

Huck, S., Normann, H.-T., and Oechssler, J. (2004a). Through trial and error to collusion*. *International Economic Review*, 45(1):205–224.

Huck, S., Normann, H.-T., and Oechssler, J. (2004b). Two are few and four are many: number effects in experimental oligopolies. *Journal of Economic Behavior & Organization*, 53(4):435–446.

Huck, S. and Oechssler, J. (2000). Through trial & error to collusion–the discrete case. *Bonn Econ Discussion Paper No. 6/2000*.

Kimbrough, S. O. and Murphy, F. H. (2009). Learning to collude tacitly on production levels by oligopolistic agents. *Computational Economics*, 33(1):47–78.

Kirman, A. P. and Vriend, N. J. (2001). Evolving market structure: An ACE model of price dispersion and loyalty. *Journal of Economic Dynamics and Control*, 25(3):459–502.

Kochanski, T. (2009). An agent based Cournot simulation with innovation: Identifying the determinants of market concentration. *MPRA Paper No. 16539*.

Leonard, D. and Nishimura, K. (1999). Nonlinear dynamics in the Cournot model without full information. *Annals of Operations Research*, 89:165–173.

Mantel, R. R. (1974). On the characterization of aggregate excess demand. *Journal of Economic Theory*, 7(3):348–353.

Nagle, T. T. and Holden, R. K. (2006). *The strategy and tactics of pricing*. Prentice Hall Englewood Cliffs, NJ, 4th edition.

Okuguchi, K. and Yamazaki, T. (2008). Global stability of unique nash equilibrium in Cournot oligopoly and rent-seeking game. *Journal of Economic Dynamics and Control*, 32(4):1204–1211.

Rand, D. (1978). Exotic phenomena in games and duopoly models. *Journal of Mathematical Economics*, 5(2):173–184.

Riechmann, T. (2002). Cournot or Walras? Agent Based Learning, Rationality, and Long Run Results in Oligopoly Games. Diskussionspapier 261, Universität Hannover, Fachbereich Wirtschaftswissenschaften.

Schipper, B. C. (2009). Imitators and optimizers in cournot oligopoly. *Journal of Economic Dynamics and Control*, 33(12):1981–1990.

Tesfatsion, L. (2006). Agent-based computational economics: A constructive approach to economic theory. *Handbook of computational economics*, 2:831–880.

Theocharis, R. (1960). On the stability of the Cournot solution on the oligopoly problem. *The review of economic studies*, 27(2):133–134.

Vega-Redondo, F. (1997). The evolution of walrasian behavior. *Econometrica*, 65(2):375–384.

Waltman, L. and Kaymak, U. (2008). Q-learning agents in a Cournot oligopoly model. *Journal of Economic Dynamics and Control*, 32(10):3275–3293.

Watkins, C. J. and Dayan, P. (1992). Q-learning. *Machine learning*, 8(3-4):279–292.

Wilensky, U. and Rand, W. (2007). Making models match: Replicating an agent-based model. *Journal of Artificial Societies and Social Simulation*, 10(4):2.

Winston, W. L. (2004). *Operations research: applications and algorithms.* Brooks Cole: Belmont, CA, 4th edition.

Zhang, T. and Brorsen, B. W. (2011). Oligopoly firms with quantity-price strategic decisions. *Journal of Economic Interaction and Coordination*, 6(2):157–170.

Abstract

Traditional oligopoly theory predicts production quantities at the Nash equilibrium level when firms have perfect information. A learning algorithm which allows firms to converge to the joint-profit maximizing outcome, even if they have only very limited information, is provided by Huck and Oechssler (2000). On the other hand, an agent-based model with a very similar learning structure results in the Cournot solution, i.e., the Nash equilibrium level (Kimbrough and Murphy, 2009). It is a-priori unclear which exact ingredient of the two models causes the strikingly different results. I identify the key difference between the models by introducing step-wise changes to the agent-based model and verifying the results in each step. The findings suggest that the key difference in the models is strict rigidity, which is absent in the model of Kimbrough and Murphy. In their model, firms rather explore than optimize in their attempt to maximize profits, involving a strong random component. This makes simultaneous downward movements very unlikely, which would be beneficial for both firms. In the model of Huck and Oechssler, firms maintain their direction of movement as long as this is beneficial for them. According to these results, inclusion of a random component in a model might lead to a drastically different outcome.

Laut klassischer Oligopoltheorie nähert sich die Gesamtproduktionsmenge der Firmen den Produktionsmengen im Nash-Gleichgewicht an, wenn Firmen perfekte Information haben. Ein Lernalgorithmus von Huck und Oechssler (2000) ermöglicht Firmen eine Gleichgewichtsproduktionsmenge in Höhe der Monopolmenge zu erreichen. Ein agentenbasiertes Modell von Kimbrough und Murphy (2009) jedoch, das auf einer sehr ähnlichen Struktur basiert, resultiert im Nash-Gleichgewicht. Welche exakte Modellkomponente für diese divergierenden Ergebnisse verantwortlich ist, ist a-priori nicht ersichtlich. Um den Hauptunterschied zwischen den beiden Modellen zu finden wird das Modell von Kimbrough und Murphy stückweise angepasst und die Resultate in jedem Schritt mit denen von Huck und Oechssler verglichen. Ich stelle fest, dass der Hauptunterschied strikte Rigididät ist, die Firmen im Modell von Kimbrough und Murphy fehlt. In deren Modell erkunden Firmen den Produktionsraum in ihrem Versuch, Profite zu maximieren. Dabei werden sie jedoch stark von einer zufälligen Komponenten beeinflusst. Dadurch werden gleichzeitige Produktionsreduktionen sehr unwahrscheinlich, die für alle Firmen vorteilhaft wären. Im Modell von Huck und Oechssler behalten Firmen die Richtung der Änderung der Produktionsmengen bei, solange dies vorteilhaft für sie ist. Entsprechend dieser Feststellung kann die Berücksichtigung einer zufälligen Komponente die Resultate eines Modells drastisch verändern.

Simon MARTIN

Education

03/2013 –
07/2014 **Mag. (equiv. MSc.)**, *University of Vienna*.
Master program Economics, Academic specialization; expected to finish in June 2014, Master thesis: Agent-based models of Cournot oligopolies, advisor: Prof. Maarten Janssen

10/2009 –
06/2010 **Pursued master courses**, *Vienna University of Economics and Business Administration*, Master program economics.

10/2009 –
06/2010 **Pursued master courses**, *Vienna University of Economics and Business Administration*, Master program Management Information Systems.

10/2006 –
06/2009 **BSc.**, *Vienna University of Economics and Business Administration*, Bachelor in Business, Economics and Social Sciences.
Major: Management Information Systems, Bachelor thesis: "Predicting a Taxpayer's Payment Behaviour"

09/2000 –
06/2005 **Final exam**, *Higher Technical Institute, focusing on computing and organization*, Leonding (Austria).
Graduated with distinction; thesis: "Spring, Struts, Hibernate - Introducing high-end J2EE frameworks"

Work experience

03/2014 -
06/2014 **Assistent in the research project "Impact of R&D subsidies"**, *Austrian Institute of Technology (AIT), Business Unit "Research, Technology & Innovation Policy"*, Vienna, Austria.
On behalf of the Austrian Federal Ministry for Traffic, Innovation and Technology (BMVIT), we analyze the impact of different research and development subsidies of Austrian government agencies onto the research output in Austria, measured by patent applications. I am responsible for modeling as well as representing the conceptual model in an agent-based simulation, building upon the simulation toolkit Mason.

12/2010 –
12/2012 **IT-Consultant**, *wedoIT-solutions GmbH*, Vienna, Austria; deployed to the Indirect Taxation Authority (ITA) in Banja Luka, Bosnia-Hercegovina.
On behalf of the Austrian Ministry of Finance and as part of a Twinning Project funded by the European Commission, I participated in a project to upgrade the IT system of the Indirect Taxation Authority (ITA) of Bosnia-Hercegovina and to contribute to further alignment with the EU acquis.
I was in charge of requirements analysis with the beneficiary, coordinating a team of software developers and software testing. Apart from that, I compiled regular reports for the EU delegation in Sarajevo.

05/2009 –
07/2009 **IT-Consultant**, *wedoIT-solutions GmbH*, Vienna, Austria; deployed to the Customs Administration of Montenegro in Podgorica.
On behalf of the Austrian Ministry of Finance and as part of a Twinning Project funded by the European Commission, I participated in a project to create a Case Management System for the Internal Audit Department of the Customs Administration of Montenegro. Also, we provided a foundation to be used for further software development projects within the administration.

01/2008 – **IT-Consultant**, *wedoIT-solutions GmbH*, Vienna, Austria.
03/2013 The company is mainly active in the area of consulting finance administrations all over the world, among them the Austrian Federal Ministry of Finance.
During this time, I gained solid knowledge in the TARIC and VAT systems in general as well as public tendering procedures in the EU and in Austria.
I supported the IT-staff in requirements analysis, testing (application of 'conformance tests' with the EU-Commission), setting up proper software architecture and consulting activities regarding the configuration and release management of software projects.

11/2006 – **Software Engineer & Architect**, *atwork information technology gmbh*, Vienna,
11/2007 Austria.
atwork is a small company that develops, among other things, the internal and external web application for the Kuratorium Wiener Pensionisten-Wohnhäuser (KWP). I supported them in that regard and drove small projects and thereby experienced the pros and cons of little firms in the IT-sector.

2004 – **Software Engineer**, *Infoniqa Informationstechnik GmbH*, Wels, Austria.
2006 I participated in the international HR-Management solution engage!, which is now used by many big and important companies, especially in the DACH-area; Gartner KG, Neumann International AG, PORR AG, REWE Austria AG, STRABAG AG, TPA-Horvath, to name only a few of them.
I gained solid knowledge both from a technical (J2EE frameworks, Apache Tomcat application server, integration with existing applications) and a functional (organizational structures, skill management, workflow) point of view.

Languages

German	**native**
English	**proficient**
Serbian / Croatian	**advanced**
Slovak	**beginner**