



CENTRE FOR DECISION RESEARCH & EXPERIMENTAL ECONOMICS



The University of
Nottingham

UNITED KINGDOM • CHINA • MALAYSIA

Discussion Paper No. 2015-10

Dávid Kopányi

May 2015

The Coexistence of Stable Equilibria under Least Squares Learning

CeDEx Discussion Paper Series
ISSN 1749 - 3293



CENTRE FOR DECISION RESEARCH & EXPERIMENTAL ECONOMICS

The Centre for Decision Research and Experimental Economics was founded in 2000, and is based in the School of Economics at the University of Nottingham.

The focus for the Centre is research into individual and strategic decision-making using a combination of theoretical and experimental methods. On the theory side, members of the Centre investigate individual choice under uncertainty, cooperative and non-cooperative game theory, as well as theories of psychology, bounded rationality and evolutionary game theory. Members of the Centre have applied experimental methods in the fields of public economics, individual choice under risk and uncertainty, strategic interaction, and the performance of auctions, markets and other economic institutions. Much of the Centre's research involves collaborative projects with researchers from other departments in the UK and overseas.

Please visit <http://www.nottingham.ac.uk/cedex> for more information about the Centre or contact

Suzanne Robey
Centre for Decision Research and Experimental Economics
School of Economics
University of Nottingham
University Park
Nottingham
NG7 2RD
Tel: +44 (0)115 95 14763
Fax: +44 (0) 115 95 14159
suzanne.robey@nottingham.ac.uk

The full list of CeDEx Discussion Papers is available at

<http://www.nottingham.ac.uk/cedex/publications/discussion-papers/index.aspx>

The Coexistence of Stable Equilibria under Least Squares Learning*

Dávid Kopányi[†]

University of Nottingham

May 28, 2015

Abstract

This paper illustrates that least squares learning may lead to suboptimal outcomes even when the estimated function perfectly fits the observations used in the regression.

We consider the Salop model with three firms and two types of consumers that face different transportation costs. Firms do not know the demand structure and they apply least squares learning to learn the demand function. In each period, firms estimate a linear perceived demand function and they play the perceived best response to the previous-period price of the other firms.

This learning rule can lead to three different outcomes: a self-sustaining equilibrium, the Nash equilibrium or an asymmetric learning-equilibrium. In this last equilibrium one firm underestimates the demand for low prices and it attracts consumers with high transportation costs only. This type of equilibrium has not been found in the literature on least squares learning before. Both the Nash equilibrium and the asymmetric learning-equilibrium are locally stable therefore the model has coexisting stable equilibria.

JEL classification: C62, C63, C72, L13

Keywords: Salop model, least squares learning, heterogeneous consumers

*I am grateful to Jan Tuinstra for his valuable comments that substantially improved the paper. I would also like to thank the participants of the 16th International Symposium on Dynamic Games and Applications (University of Amsterdam), the Brown Bag Seminar (University of Nottingham), the 2015 Annual Conference of the Royal Economic Society (University of Manchester) and the 20th Annual Workshop on the Economic Science with Heterogeneous Interacting Agents (SKEMA Business School, Sophia Antipolis) for useful suggestions and feedback.

[†]email: david.kopanyi@nottingham.ac.uk

1 Introduction

There are many situations where firms do not fully know the demand conditions on the market. For example, when a firm is entering a new market or when a new product is introduced, firms need to learn over time how the demand for their good depends on the price they set and on their competitors' prices. Similarly, changes in market regulation can also lead to demand uncertainty. In such situations firms need to learn about demand conditions to find the optimal price or production level.

A natural way to learn about unknown demand conditions is to gather market data and estimate the demand function. Then this estimated function can form the basis for determining prices or quantities. This process can be modeled with *least squares learning* (LSL). LSL consists of two parts: estimation and a decision rule. In the estimation part agents estimate an unknown relationship between variables with OLS regression and then they use the estimated function to determine their optimal action. If the functional form in the estimation is correctly specified, i.e. agents have observations about all the relevant variables and these variables are included in the regression in the correct functional form, then LSL leads to the rational expectations equilibrium, see Bray (1982), Marcet and Sargent (1989) and Evans and Honkapohja (2001) for example. However, since the true relationship is unknown, it might occur that agents use a misspecified functional form in the regression. The effect of such misspecification was analyzed in the Industrial Organization literature, see Kirman (1975, 1983), Gates et al. (1977), Brousseau and Kirman (1992), Kopányi (2013) and Anufriev et al. (2013). Misspecified LSL leads to a so-called self-sustaining equilibrium in which firms do not learn the true demand function correctly but they end up in a situation where 1. they are choosing the profit-maximizing price or quantity subject to their estimated function and 2. in the equilibrium their estimation is correct in the sense that the price or demand they *expect* to get (based on their estimation) coincides with the actual realization.¹ This kind of outcome is unrelated to any benchmark outcomes of the standard static model under complete knowledge about the demand structure, such as the Nash equilibrium or the collusive outcome.²

In this paper we take an intermediate step between the aforementioned branches of the literature. We take a slight departure from the correctly specified version of LSL by assuming that firms use all the relevant variables in the regression (i.e. the prices set by their competitors) and that the functional

¹Thus, their estimation is correct in the equilibrium point but not out of equilibrium.

²We must note that the result that misspecified learning can lead to different outcomes than a correctly specified learning rule holds not only for least squares learning. Bayesian learning leads to similar results depending on whether the likelihood function is correctly specified or not, see Blume et al. (1982) and the references therein. For market models with misspecified Bayesian learning see Blume and Easley (1982) and Schinkel et al. (2002) for example.

form they use is correctly specified locally but not globally.³ This situation is of particular interest as the correctly specified and misspecified versions of LSL lead to substantially different outcomes and it is unclear whether a small misspecification in the learning method, which can easily occur and is hard to detect, is enough to prevent convergence to the Nash equilibrium. As our results show, LSL can lead both to the Nash equilibrium and to a self-sustaining equilibrium. Moreover, a third kind of outcome, which was not present in previous models, can also be reached: the asymmetric learning-equilibrium. In this equilibrium some firms charge higher prices than others and they focus only on a small part of the market. A remarkable feature of this outcome is that even though firms observe a perfect fit, the outcome can be suboptimal for some firms in the sense that they could make a higher profit by focusing on the whole market. However, in some situations it can be optimal to focus on only one part of the market.

As a framework of the analysis we consider a modified version of the circular road model introduced by Salop (1979). Three firms produce a homogeneous good. Firms are located along a circular road, in equidistant locations. Consumers are uniformly distributed along the circle. When a consumer wants to buy the good, it needs to visit one of the firms. Transportation is costly, consumers face a fixed transportation cost per distance unit. Thus, the total cost of buying the good from a specific firm is given by the sum of the price the firm asks and the transportation costs. Demand is inelastic, each consumer is assumed to buy exactly one unit of the good, at the lowest possible total cost. We introduce heterogeneity on the consumer side. There are two types of consumers, one type faces low transportation cost while the other type faces a high one.

Firms do not know the market structure and they use LSL to learn the demand function they face. The true demand function is piecewise linear but firms approximate it with a linear function. Hence the approximation can be locally correct but globally incorrect as a firm can get a correct approximation for at most one of the linear parts of the true demand function. In this paper we investigate which outcomes LSL can lead to in this situation. We analytically show that the model has three kinds of equilibria. When firms use all past observations in the estimation, LSL typically leads to a *self-sustaining equilibrium*. In this equilibrium firms choose the price that maximizes their expected profit subject to their beliefs about demand conditions and their beliefs are correct in equilibrium but they are incorrect out of equilibrium. On the other hand, when not all but only the most recent observations are used in the estimation, firms reach either the symmetric *Nash equilibrium* or the *asymmetric learning-equilibrium*. In this asymmetric learning-equilibrium two firms charge a low price and the third one asks a high price. The high-price firm attracts the high-type consumers only whereas the other two firms serve both consumer types. The

³As we will explain later, they approximate a piecewise linear function with a linear one.

intuition behind this equilibrium is that the high-price firm does not attract low-type consumers, therefore it underestimates the demand at low prices and it does not perceive it profitable to charge a lower price. We analytically investigate which conditions determine the outcome of the learning process and we run numerical simulations to evaluate how frequently the different outcomes are reached.

Least squares learning was applied in market competition in other papers as well. See Anufriev et al. (2013) for an overview of the literature on misspecified LSL. Our results are in line with the findings of this literature when firms use all past observations in the regression. In the Nash equilibrium and in the asymmetric learning-equilibrium, perceived demand functions are correctly specified in the neighborhood of the equilibrium price. This makes these outcomes more robust than the self-sustaining equilibria in Brousseau and Kirman (1992) and Anufriev et al. (2013) in the sense that in case of an SSE a firm would discover that its perceived demand function is misspecified by choosing a slightly different price. This is not the case for the Nash equilibrium and the asymmetric learning-equilibrium.⁴

Tuinstra (2004) takes a similar approach as we do in the sense that he considers a perceived demand function that is locally correct but globally incorrect. In his paper, the perceived demand function is the linear approximation of the true nonlinear demand function at the current price vector (i.e. the perceived demand function matches the function value and the slope of the true demand function at the current price).⁵ Thus, the approximation is correct at the equilibrium point only, whereas it is correct in a *neighborhood* of an equilibrium in our model (in case of the Nash equilibrium and the asymmetric learning-equilibrium). Another important difference is that firms focus only on their own price effect in the approximation in Tuinstra's paper while they take into account the prices of other firms as well in our model.

The paper is structured as follows. The circular road model is discussed in Section 2. In Section 3 we discuss least squares learning and we derive the equilibria of the model. We analyze the stability of the equilibria as well. Simulation results are reported in Section 4. Section 5 concludes. Proofs are presented in Appendix A while robustness checks are reported in Appendix B.

2 The circular road model

The circular road model, one of the baseline models of horizontal product differentiation, was introduced by Salop (1979). In this section we first review a simplified version of the model that is relevant for our

⁴Note, however, that this difference is due to the different informational structure of the models. In Brousseau and Kirman (1992) and Anufriev et al. (2013) firms can observe their own actions only, whereas they have full information about the actions in our model.

⁵One of the equilibrium concepts in Silvestre (1977) is based on similar conditions.

analysis and then we introduce heterogeneity on the consumer side.

2.1 Homogeneous consumers

Consider the market for a homogeneous good that is produced by three firms. Firms simultaneously and independently set the price of the good. Production costs are given by the same function for each firm: $C_i(q_i) = cq_i$ for each firm i , where q_i is the production level of firm i and $c > 0$ is a parameter. Firms are located along a circular road, in equal distance from each other. Consumers are uniformly distributed along the circle, their mass (or equivalently the circumference of the circle) is normalized to 1.

Consumers need to visit one of the firms to purchase the good. They move along the circular road, facing a transportation cost s per distance unit. If the minimal distance between firm i and a given consumer is x , then the consumer's total cost for buying the good from firm i is $p_i + sx$, where p_i is the price charged by firm i and sx is the total transportation cost.⁶ Demand is inelastic: each consumer buys exactly one unit of the good. Furthermore, consumers are assumed to buy the good at the lowest possible cost, thus from the firm for which the sum of the price and the total transportation cost is the lowest. An alternative interpretation of the model is that the circle represents the product space and the location of consumers determines their preferences for the different products. Consumers choose a product based on the prices and on the distances from their ideal product, which corresponds to their location.

Firms can drive out each other from the market by choosing a sufficiently low price. When a consumer at the location of a firm is better-off by visiting another firm, then the first firm will not attract any consumers. This setup results in demand functions that are discontinuous and which consist of piecewise linear parts. This one-shot game has a unique symmetric Nash equilibrium, in which each firm charges the price $p = c + \frac{s}{3}$, see Tirole (1988), p. 283 for the proof. Having discussed how the basic model works, let us introduce heterogeneity on the consumer side.

2.2 Heterogeneous consumers

Let us consider the same market structure as before but suppose that there are two types of consumers. The types differ with respect to the transportation cost they face: low-type consumers face a unit cost of s while high-type consumers pay a unit cost of S , where $s < S$. The amount of consumers of each type is normalized to 1, both types are assumed to be uniformly distributed along the circular road. Firms cannot distinguish the two types, they cannot price discriminate between different consumers.⁷

⁶It is assumed that firms cannot price discriminate so they cannot charge different prices to consumers from different locations.

⁷We also assume that resale is not possible.




Figure 1: Demand function of firm i for $p_j = 2$ and $p_k = 2.2$. Other parameters: $s = 1$, $S = 2$ and $c = 1$.

Similarly to the case with homogeneous consumers, firms can drive each other out of the market by choosing a sufficiently low price. Moreover, firms can also be driven out of one part of the market only: It can occur that a firm attracts high-type consumers but not low-type ones. Consider for example the situation $p_j + \frac{1}{3}s < p_i < p_j + \frac{1}{3}S$. In this case the low-type consumer that is located at the position of firm i buys from firm j rather than from firm i . Consequently, firm i does not attract low-type consumers. On the other hand, the high-type consumer at the location of firm i prefers firm i to firm j . Thus, in the given situation, firm j drives firm i out of the market for low-type consumers but not for high-type consumers.

Demand functions are discontinuous and consist of piecewise linear parts again. There are more parts than under homogeneous consumers since firms can be driven out of multiple subparts of market in this case. We do not report the exact formula for the demand function here as it is not important to know it for understanding the results of the paper. The relevant linear parts of the demand function are derived in Appendix A. Figure 1 illustrates the demand and profit functions of firm i when the other two firms do not drive each other out of the market for either consumer type. We can see that the demand function indeed consists of linear parts. There are 7 linear parts, they correspond to the following cases (as p_i increases): 1. firm i serves the whole market; 2. low-type consumers are served by firm i only, high-type consumers are served by firms i and j ; 3. low-type consumers are served by firm i only, high-type consumers are served by all 3 firms; 4. low-type consumers are served by firms i and j , high-type consumers are served by all 3 firms; 5. both consumer types are served by all three firms; 6. firm i serves high-type consumers only, the other two firms serve both consumer types; and 7. firm i is completely driven out of the market. We can see from the profit function that the profit-maximizing decision of firm i in the given situation is to drive

the other two firms out of the market for the low-type consumers but not for the high-type, as the profit maximum is reached in the third case.

The model with heterogeneous consumers has a unique Nash equilibrium in pure strategies. Proposition 2.1 specifies the equilibrium price. The proof of the proposition is presented in Appendix A.

Proposition 2.1 *The Salop model with three firms and two types of consumers has a unique Nash equilibrium in pure strategies. This equilibrium is symmetric, with all three firms charging the price $p_N = \frac{2Ss}{3(S+s)} + c$ and serving both consumer types.*

Note that the proposition rules out the existence of asymmetric Nash equilibria. The Nash equilibrium price is increasing in both s and S . The intuition behind this result is the following. When transportation costs are higher, it is harder for firms to attract consumers that are located farther away from them (or equivalently, it is more costly for consumers to visit firms that are farther away from them). This reduces competition, firms gain more market power and the equilibrium price increases consequently.

It can be seen that $\frac{\partial p_N}{\partial S} < \frac{\partial p_N}{\partial s}$, that is s has a larger impact on the equilibrium price than S does. To understand this result, note the following. When a transportation cost increases, firms have an incentive to increase their price since they get more market power in the given market segment. When a firm increases its price, it will lose some low-type as well as high-type consumers. Since low-type consumers are more mobile, the firm will lose more low-type consumers. Thus, it is more favorable for firms when the transportation cost of low-type consumers increases since this makes low-type consumers less mobile, resulting in a lower decrease in demand after a price increase. Thus, the equilibrium price increases more when s increases.

After analyzing the static model under full information, we now turn to a dynamic model in which firms do not know the market specification and they try to learn the demand for their good using market observations.

3 Market dynamics under learning

When firms do not know the market structure, they need to learn the demand function to find the optimal action. When firms apply least squares learning, they approximate the true demand function with a *perceived* demand function and they estimate the unknown parameters of it using past observations about prices and production levels.

We assume that the only information the firms have about the market is that there are three firms. Thus, they do not know either about the circular-road structure of the market or about consumer heterogeneity. Firms are competing with each other on the same market over time and they can observe the prices

charged by their competitors and the corresponding demand for their own good (but not those of their competitors). Thus, firms gather information about the market over time and they use this information to learn about the demand for their product.

In the following subsection we specify the learning method the firms use and then we discuss the equilibria of the model under learning.

3.1 Least squares learning

Firms approximate the demand for their product with a linear function. The *perceived* demand function of firm i is given by

$$D_i^P(p) = a_i - b_{ii}p_i + b_{ij}p_j + b_{ik}p_k + \varepsilon_i, \quad (1)$$

where a_i denotes the demand intercept, b_{ix} denotes the effect of firm x 's price on the demand for firm i 's product ($x = i, j, k$) and ε_i is a random variable with mean 0. Parameters a_i and b_{ix} are estimated with OLS regression using observations about past prices and own-production levels.

Firms might not want to use all past observations for the estimation therefore we need to make a distinction between a firm's observations and information set. *Observations* of firm i consist of the prices of all three firms and the demand firm i faces for all past periods whereas the *information set* contains only those observations that are used in the regression.⁸ The rationale behind not using all observations in the regression is that older observations might carry less information about current demand conditions than more recent ones, especially when there is a structural break in the data. Even though demand conditions are fixed in the model we consider, not using all past observations, as we will see, has important consequences for the properties of LSL.

Let us suppose that firms use the last τ observations in the regression. Then parameter estimates for firm i are given by the standard OLS formula

$$\beta_i = (X'_{i,\tau} X_{i,\tau})^{-1} X'_{i,\tau} y_{i,\tau}, \quad (2)$$

where $\beta_i = (a_i, b_{ii}, b_{ij}, b_{ik})'$ is the 4×1 vector of parameter estimates⁹, $X_{i,\tau}$ is the $\tau \times 4$ matrix of explanatory variables (constant and price observations for the last τ periods) and $y_{i,\tau}$ is the $\tau \times 1$ vector of the last τ demand observations of firm i (dependent variable).¹⁰

⁸Note that we use the term information set in its econometric sense and not in its game theoretical sense.

⁹To simplify notation, we denote the unknown parameters of the perceived demand function as well as the corresponding parameter estimates by the same symbol. This should not be confusing as we will work only with parameter estimates from now on.

¹⁰Similar formulas apply when firms use all past observations. The only difference is that X and y then contain the prices and the corresponding demand for all past periods.

Given the parameter estimates of the perceived demand function, firm i maximizes its one-period expected profit $E\pi_i^P(p) = (p_i - c)ED_i^P(p)$. This leads to the following best-response price:

$$p_i^{\text{BR}} = \frac{a_i + b_{ij}p_j + b_{ik}p_k}{2b_{ii}} + \frac{c}{2}. \quad (3)$$

Let us now discuss timing. At the end of period t firms have observations about all t periods. Parameter estimates are obtained by (2). In order to stress that parameter estimates are changing over time, we will denote the parameter estimates at the end of period t as $a_{i,t}$, $b_{ii,t}$, $b_{ij,t}$ and $b_{ik,t}$. Since firms are determining their prices simultaneously, they can play the best response only against the *expectations* they have about the prices of other firms. Thus, we have to replace p_j with $p_{j,t+1}^e$ and p_k with $p_{k,t+1}^e$ in (3), stressing again the dependence on time. We assume that firms form naive expectations, meaning that they expect other firms to charge the same price as in the previous period: $p_{j,t+1}^e = p_{j,t}$ and $p_{k,t+1}^e = p_{k,t}$. This leads to the following pricing formula for period $t + 1$:

$$p_{i,t+1} = \frac{a_{i,t} + b_{ij,t}p_{j,t} + b_{ik,t}p_{k,t}}{2b_{ii,t}} + \frac{c}{2}. \quad (4)$$

Note that profit maximization requires $b_{ii,t} > 0$, that is the perceived own-price effect must be negative. Since the perceived demand functions the firms use are not correctly specified, the parameter estimate for $b_{ii,t}$ might become negative. In this case, (4) does not give the perceived profit-maximizing price. Also note that when (4) yields a price that is lower than the marginal cost, the firm would make a negative profit (provided that it faces a positive demand). Thus, (4) is not applicable in this case either. In order to overcome these possible issues with LSL, we augment the method with the following rule.

Random price rule: When $b_{ii,t} \leq 0$ or (4) yields $p_{i,t+1} < c$, then firm i chooses a price randomly from the uniform distribution on a predefined interval I .

Interval I is specified in Section 4. We need to impose additional rules to overcome some numerical issues that may occur when firms do not use all observations in the regression. When prices start to settle down at a given value, there might not be enough dispersion in the observations and matrix $X_{i,\tau}$ can be close to being singular, resulting in imprecise parameter estimates. This can lead to extremely high prices for some periods. Since it should be clear for firms that large unexpected price changes result from the aforementioned issue, it is reasonable to assume that firms do not follow pricing rule (4) in this case, they rather keep their price unchanged. This leads to the following rule.

No jump rule: If (4) yields a price that is at least K times higher than the price of firm i in the previous

period, then the firm will keep its price unchanged and charge the same price as in the previous period.¹¹

When there is not enough dispersion in the price observations, matrix $X_{i,\tau}$ can become singular, making the estimation impossible. We assume that firms keep their price unchanged in this case.

Impossible estimation rule: When (2) is not applicable due to the singularity of $X_{i,\tau}$, then firm i will keep its price unchanged and charge the same price as in the previous period.

It might seem as if the last two rules lead to an artificial stability in the model as we require firms to use the same price as in the previous period but actually these rules exclude an artificial instability. Note that the estimation problem occurs only when the process has almost converged. Thus, firms observe that prices have settled down around some values and then the new parameter estimates lead to an unexpectedly large price. First of all firms might be reluctant to make such a big price change, secondly after observing the time series of prices it should be clear that this sudden price change comes from a numerical issue, therefore it is better not to change the price. Concerning the impossible estimation rule, when parameter estimates cannot be obtained, then firms either choose a price randomly or they fix the price as we suggest.¹² Keeping the price unchanged is a more reasonable solution to this issue. In Section 4 we discuss how often the different rules are used in the simulations.

Let us now turn to the steady states of the process.

3.2 Equilibria under least squares learning

The system is in a steady state when neither the parameter estimates of the perceived demand functions nor the prices change. It must hold for any steady state that the true and (expected) perceived demand functions coincide for each firm at the given price vector p^* , that is $D_i(p^*) = ED_i^P(p^*)$ for $i = 1, 2, 3$. To see this, note the following. When $D_i(p^*) = ED_i^P(p^*)$, the perceived demand function perfectly approximates the true demand function for the given price vector as the corresponding estimation error is 0. Since the parameter estimates of the perceived demand function are obtained by minimizing the sum of squared errors, this implies that the parameter estimates do not change in this case.

¹¹Alternatively, we could impose an upper bound on price changes as Weddepohl (1995). In that case firms would choose the highest possible price if (4) resulted in a too large price jump. Since large price jumps are associated with imprecise parameter estimates in the model we consider, it makes more sense not to change the price at all.

¹²Note that a rule like the no jump rule or the impossible estimation rule is essential for having convergence in a model that is not subject to noise (e.g. demand shock) when firms do not use all past observations in the regression. If the process converged to a certain value, then estimation would not be possible since each observation would perfectly correspond to the steady state. Thus, it needs to be specified what happens when parameter estimates cannot be calculated.

The same condition characterizes the self-sustaining equilibria in Brousseau and Kirman (1992) and in Anufriev et al. (2013). Thus, the steady states of the model with least squares learning are self-sustaining equilibria: firms play the best response subject to their beliefs about demand conditions (i.e. the perceived demand functions) and about the prices of the other firms, and these beliefs are correct at the equilibrium price vector. Self-sustaining equilibria can be formally defined as follows.

Definition 3.1 *Price vector $p^* = (p_1^*, p_2^*, p_3^*)$ and the parameter estimates $\{a_i^*, b_{ii}^*, b_{ij}^*, b_{ik}^*\}$ ($i, j, k = 1, 2, 3; i \neq j \neq k$) constitute a self-sustaining equilibrium if the following conditions hold for each firm i :*

$$p_i^* = \frac{a_i^* + b_{ij}^* p_j^* + b_{ik}^* p_k^*}{2b_{ii}^*} + \frac{c}{2}, \quad (5)$$

$$ED_i^P(p^*) = D_i(p^*). \quad (6)$$


Condition (5) shows that firms play the best response subject to their beliefs and (6) means that beliefs are confirmed in equilibrium as the actual demand is the same as the demand the firm expects to get, and the prices of the competitors are also as expected.

It can be seen from the definition that there are many different self-sustaining equilibria, thus the model has multiple steady states. Proposition 3.2 specifies which price vectors can form a self-sustaining equilibrium.


Proposition 3.2 *For any price vector $p = (p_1, p_2, p_3)$ satisfying the conditions $p_i > c$ and $D_i(p) > 0$ for $i = 1, 2, 3$, there exist values of $\{a_i, b_{ii}, b_{ij}, b_{ik}\}$ ($i, j, k = 1, 2, 3; i \neq j \neq k$) such that the model is in a self-sustaining equilibrium.*

Thus, prices exceed the marginal cost and each firm faces a positive demand in a self-sustaining equilibrium. Note that the condition $D_i(p) > 0$ implies that none of the firms can be driven out of the market for *both* types of consumers. But it is not required that each firm should attract both consumer types. In the above result, we did not take into account that $\{a_i, b_{ii}, b_{ij}, b_{ik}\}$ are not freely chosen but they result from estimation. Therefore not all the price vectors that satisfy the conditions of Proposition 3.2 can necessarily be *reached*, despite the fact that we can find parameter values for which they constitute a self-sustaining equilibrium.

Since perceived demand functions are linear while the true demand functions are piecewise linear, firms cannot fully learn the true demand conditions: They can correctly learn the parameters of at most one linear part. Note that condition (6) is required to hold at the equilibrium point only, thus firms need not learn in general any linear part correctly. Panel *a* of Figure 2 illustrates the true and the perceived



(a) Demand functions in an SSE



(b) Profit functions in an SSE


Figure 2: Demand and profit functions of firm i in a self-sustaining equilibrium. Parameters: $s = 1$, $S = 5$ and $c = 1$. Equilibrium prices: $p_i^* = 2.0398$, $p_j^* = 2.0264$ and $p_k^* = 2.2083$.

demand functions of a firm in a typical self-sustaining equilibrium. The two functions cross each other in a single point thus the firm does not learn any linear part of the true demand function correctly. Panel b depicts the true and the perceived profit functions. The figure shows that in the SSE firm i maximizes its perceived profit but the price it chooses does not yield the true profit maximum.


Even though it is not the case for a general SSE, there are self-sustaining equilibria in which firms correctly learn the part of the true demand function on which they operate. Proposition 3.3 specifies these equilibria.

Proposition 3.3 *The model with least squares learning has two self-sustaining equilibria in which firms correctly learn that linear part of the true demand function on which they operate. The Nash equilibrium of the static game is always such an equilibrium of the learning process. When $\frac{S}{s} \geq \Sigma_1 = \frac{7+\sqrt{89}}{4} \approx 4.1085$, there also exists another equilibrium in which two firms charge $p_L = \frac{11Ss}{12S+15s} + c$ and the third firm chooses $p_H = \frac{2S^2+8Ss}{12S+15s} + c$. We refer to this equilibrium as asymmetric learning-equilibrium (ALE).*


Figure 3 illustrates the demand and profit functions in the Nash equilibrium and in the asymmetric learning-equilibrium. Panels a , c and e confirm that in both equilibria firms correctly approximate the linear part of the true demand function on which they operate. Panel b shows that the true profit maximum coincides with the maximum of the perceived profit function of firms in the Nash equilibrium. The same holds for the low-price firms in the ALE (see panel d). Note, however, that the perceived profit maximum does not correspond to the true profit maximum for the high-price firm (panel f), it reaches a local profit




(a) Demand functions (Nash equilibrium)




(b) Profit functions (Nash equilibrium)




(c) Demand functions of low-price firms (ALE)



(d) Profit functions of low-price firms (ALE)



(e) Demand functions of the high-price firm (ALE)



(f) Profit functions of the high-price firm (ALE)

Figure 3: Demand and profit functions in the Nash equilibrium and in the asymmetric learning-equilibrium. Parameters: $s = 1$, $S = 5$ and $c = 1$.

maximum only.¹³ This is why the ALE is not a Nash equilibrium of the game under known demand. As panel *e* shows, the high-price firm underestimates the demand for lower prices and thus it does not perceive it more profitable to charge a lower price, even though it would yield a higher profit. The intuition behind this result is the following. If a firm asks too high prices, it will not attract low-type consumers. Since firms learn about demand conditions solely from market observations, the firm in question gets information about one part of the market only and that is why it underestimates demand for lower prices.

The following proposition compares the Nash equilibrium and the ALE in terms of prices and profits.

Proposition 3.4 *The following relationships hold between prices and profits in the Nash equilibrium and the ALE:*

- $p_N < p_L < p_H$,
- $\pi_H < \pi_N < \pi_L$ for $\frac{7+\sqrt{89}}{4} < \frac{S}{s} < 8.91$,
- $\pi_N < \pi_H < \pi_L$ for $8.91 < \frac{S}{s} < \frac{89+11\sqrt{73}}{8} \approx 22.87$,
- $\pi_N < \pi_L < \pi_H$ for $\frac{89+11\sqrt{73}}{8} < \frac{S}{s}$,
- $3\pi_N < 2\pi_L + \pi_H$.

The finding about prices is in line with the fact that prices are strategic complements in the model: the high-price firm charges a higher price than in the Nash equilibrium and this gives an incentive for the other two firms to increase their price. That is why $p_L > p_N$. Concerning profits, low-price firms always earn a higher profit than in the Nash equilibrium while the high-price firm may earn a lower as well as a higher profit. Note that when $\pi_H < \pi_N$, the high-price firm reaches a suboptimal outcome even though the perceived demand function perfectly fits the true demand function locally. For intermediate values of $\frac{S}{s}$, all three firms are better-off than in the Nash equilibrium thus the high-price firm is better-off by not attracting low-type consumers. For high values of $\frac{S}{s}$, the high-price firm earns the highest profit. In this case the high-price firm still underestimates the demand for low prices but the perceived profit maximum coincides with the true profit maximum. On the other hand, low-price firms perceive a relatively high slope and they underestimate the demand for high prices. Their perceived profit maximum does not coincide with the true profit maximum as it would be more profitable to charge a higher price. Finally, the total profit of the three firms is always higher in the ALE than in the Nash equilibrium.

¹³The ALE can be viewed as a local Nash equilibrium since the low-price firms reach their global profit maximum while the high-price firm is in a local profit maximum only. See Bonanno and Zeeman (1985) and Bonanno (1988) for more details about this concept.

Since prices are higher in the ALE than in the Nash equilibrium, consumers are worse-off. Moreover, welfare (measured as total surplus) is lower. Note that for comparing the welfare in the two outcomes, it is enough to focus on transportation costs only. The reason for this is the following. The surplus of a consumer can be measured as the net utility of consuming the good: $v - p - sx$ (or $v - p - Sx$), where $v > 0$ is the positive utility from consumption while $p + sx$ (or $p + Sx$) is the total cost of purchasing the good.¹⁴ Note that the price p is simply a transfer between the consumer and the firm, therefore it does not have a direct effect on welfare. Also note that total production is the same in the Nash equilibrium and in the ALE. Since the marginal cost of production is constant and equal for the firms, the difference in individual production levels does not contribute to welfare differences. Thus, from a welfare perspective, only transportation costs matter. Transportation costs are higher in the ALE than in the Nash equilibrium for two reasons. First, low-type consumers go to the low-price firms only, thus some of these consumers need to travel more compared to the Nash equilibrium. Second, the high-type consumers that are indifferent between the high-price firm and one of the low-price firms, lie closer to the high-price firm than under a symmetric situation (as in the Nash equilibrium). Therefore, those high-type consumers that visit the low-price firm but would visit the other firm in a symmetric situation, travel more than in the Nash equilibrium. Thus, even though total profits are higher, welfare is lower in the ALE.

3.3 Stability of equilibria

As we have seen in the previous section, the model with least squares learning has three types of equilibria: a general self-sustaining equilibrium, the Nash equilibrium and the asymmetric learning-equilibrium. Next we will investigate which equilibria can be reached and which factors determine which of the equilibria is reached. It turns out that a special property of the information set plays a crucial role in this. Before defining this property, note that different price vectors may correspond to different demand conditions. For example, firm i may serve both types of consumers for one price vector whereas it might serve high-type consumers only for another price vector. These price observations carry information about different structural parameters as they lie on different linear parts of the true demand function. We call price vectors in the information set of firms *aligned* when each firm serves the same consumer type(s) for each price vector. We distinguish two kinds of aligned price vectors. When all three firms serve both consumer types, we speak about *symmetrically aligned* prices. When two of the firms serve both consumer types while the third one attracts high-type consumers only, we speak about *asymmetrically aligned* prices.¹⁵ We define

¹⁴Remember that each consumer is assumed to buy the good. This implies that v is assumed to be sufficiently large.

¹⁵Note that prices could be aligned in other ways as well. For example, we could consider the case when one firm attracts both types of consumers while the other two firms attract high-type consumers only. We do not consider other possibilities

these concepts formally as follows.

Definition 3.5 *A set of price vectors $P \subseteq \mathbb{R}_+^3$ is called symmetrically aligned when all three firms attract both types of consumers for all $p \in P$:*

$$|p_i - p_j| < \frac{s}{3} \quad \forall i, j = 1, 2, 3.$$

A set of price vectors $P \subseteq \mathbb{R}_+^3$ is called asymmetrically aligned when firms i and j attract both types of consumers while firm k attracts only the high-type consumers for all $p \in P$:

$$|p_i - p_j| < \frac{s}{3} \\ \min\{p_i, p_j\} + \frac{s}{3} < p_k < \min\{p_i, p_j\} + \frac{S}{3}.$$

A set of price vectors $P \subseteq \mathbb{R}^3$ is called not aligned when it is neither symmetrically, nor asymmetrically aligned.

The condition $|p_i - p_j| < \frac{s}{3}$ ensures that firms i and j do not drive each other out of the market for either consumer type. The condition $\min\{p_i, p_j\} + \frac{s}{3} < p_k < \min\{p_i, p_j\} + \frac{S}{3}$ means that firm k is driven out of the market for low-type consumers but not for the high-type ones.

When prices are aligned, then the corresponding demand observations are consistent in the sense that they lie on the same linear part of the demand function. That is, observations carry information about the same linear demand parameters and consequently firms correctly learn the parameters that characterize the linear part of the true demand function on which they operate.

Since firms play the best response to the prices of the other firms, subject to their perceived demand function, it is important to analyze the conditions under which a set of aligned price observations remains aligned after updating the set with the best-response prices. Lemma 3.6 summarizes these conditions.

Lemma 3.6 *When price observations are symmetrically aligned, then updating the information set with the best-response prices always results in symmetrically aligned price observations again.*

When price observations are asymmetrically aligned, there are three possibilities.

1. *For $\frac{S}{s} < \Sigma_1$ price observations will not be asymmetrically aligned after updating the information set with the best-response prices sufficiently many times.*
2. *For $\frac{S}{s} \in [\Sigma_1, \Sigma_2)$ with $\Sigma_2 = 2 + \sqrt{6} \approx 4.4495$, the updated price observations will be asymmetrically aligned if the following condition holds for the most recent price observation p :*

$$\frac{s}{3} \left[2 \left(\frac{S}{s} \right)^2 - 7 \frac{S}{s} - 4 \right] + \left(1 + \frac{S}{s} \right) |p_i - p_j| + \min\{p_i, p_j\} \geq p_k.$$

because they are not relevant for the equilibria of the learning process, as we have seen.

3. For $\frac{S}{s} \geq \Sigma_2$, price observations always remain asymmetrically aligned after updating the information set with the best-response prices.

According to this lemma, when the information set is symmetrically aligned, then it always remains symmetrically aligned. Thus, firms will learn the true parameters of the corresponding linear part. As Proposition 3.3 shows, the only equilibrium that firms may reach in this situation is the Nash equilibrium. Concerning asymmetrically aligned observations, Lemma 3.6 says that when $\frac{S}{s}$ is not high enough, the information set will not be asymmetrically aligned eventually even if firms start with an asymmetrically aligned information set. So in this case the possible steady states of the model are a general SSE and the Nash equilibrium. For intermediate values of $\frac{S}{s}$, an extra condition is needed for ensuring that the updated information set remains asymmetrically aligned. Thus, all three steady states may exist for these values of $\frac{S}{s}$. On the other hand, an asymmetrically aligned information set always remains asymmetrically aligned by updating it with the best response prices when $\frac{S}{s}$ is high enough. Thus, the only equilibrium in this case is the asymmetric learning-equilibrium. When price observations are not aligned, then firms cannot learn the true parameters of the linear part on which they operate, consequently the only kind of steady state in the given situation is a general self-sustaining equilibrium.

Note that these results concern existence only, under specific conditions. We have not analyzed the stability of these equilibria yet. Proposition 3.7 summarizes the dynamical properties of the steady states.

Proposition 3.7 *Both the Nash equilibrium and the asymmetric learning-equilibrium are locally stable equilibria of the model with least squares learning.*

According to the proposition, firms will reach the Nash equilibrium when initial prices are close to the Nash equilibrium price. A similar result holds for the asymmetric learning-equilibrium. Combining these considerations with Lemma 3.6, we can conclude that the model has coexisting locally stable steady states when $\frac{S}{s}$ is sufficiently high. Note that Proposition 3.7 does not cover the stability of general self-sustaining equilibria. Brousseau and Kirman (1992) show that firms do not converge to a self-sustaining equilibrium in general. The process slows down only because the weight of a new observation decreases when firms use all observations in the estimation.

Taking into account the above theoretical results, we summarize the steady states of the model in Table 1. When firms use all observations in the estimation, then all three equilibria can occur. More specifically, when initial observations are symmetrically aligned, firms converge to the Nash equilibrium. When initial observations are asymmetrically aligned and $\frac{S}{s}$ is sufficiently high, then firms reach the asymmetric learning-equilibrium. When initial observations are not aligned or if they are asymmetrically aligned but $\frac{S}{s}$ is not high enough, then firms move towards a self-sustaining equilibrium.

initial observations	information set	
	all observations	last τ observations
symm. aligned	Nash	Nash
$\frac{S}{s} \leq \Sigma_1$	SSE	Nash
asymm. aligned $\Sigma_1 < \frac{S}{s} < \Sigma_2$	ALE / SSE	ALE / Nash
$\Sigma_2 \leq \frac{S}{s}$	ALE	ALE
not aligned	SSE	ALE / Nash

Table 1: The steady states of the model with least squares learning for different types of initial observations and different number of observations in the information set.

When only the last τ observations are used in the regression, then SSE cannot be reached. If there is convergence in the model, then the information set must become either symmetrically or asymmetrically aligned: if this was not the case, then there must be jumps between the different linear parts of the demand function, contradicting convergence. Thus, the Nash equilibrium and the ALE can be reached more often in this case for the following reason: an information set which is not aligned might become symmetrically or asymmetrically aligned as old observations drop out of the information set at some point. Thus, when only the most recent observations are used in the regression, the system can converge either to the Nash equilibrium or to the ALE, or there is no convergence at all. Since both the Nash equilibrium and the ALE are locally stable, we expect that observations will not jump between the different linear parts of the demand function and therefore one of the equilibria will be reached eventually. If this conjecture does not hold, then the process does not converge at all as observations keep on jumping between the different linear parts of the true demand function.

In the next section we run computer simulations to investigate whether firms indeed reach the Nash equilibrium or the ALE when only the most recent observations are used. We also investigate how often the different outcomes are reached.

4 Simulation results

We run simulations with 1000 different initializations. Each initialization runs until the maximal price change is smaller than the threshold value of 10^{-8} , i.e. $\max_i |p_{i,t} - p_{i,t-1}| \leq 10^{-8}$, or until period 1000 is reached. We fix the market parameters at $c = 1$ and $s = 1$, and we vary the value of S . Based on the theoretical results we consider 6 different values, Table 2 summarizes them as well as the corresponding

	S					
	2	4	4.2	4.35	4.5	10
p_N	1.4444	1.5333	1.5385	1.5421	1.5455	1.6061
p_L	-	-	1.7064	1.7121	1.7174	1.8148
p_H	-	-	2.0532	2.0810	2.1087	3.0741

Table 2: The Nash equilibrium price and prices in the asymmetric learning-equilibrium for different values of S . Other parameters: $s = 1$ and $c = 1$.

prices in the Nash equilibrium and in the asymmetric learning-equilibrium. The chance of reaching the Nash equilibrium and the ALE is different across these values. For $S = 2$ and $S = 4$ the ALE does not exist as $S < \Sigma_1 \approx 4.1085$. For $S = 4.2$ and $S = 4.35$ the ALE exists but an asymmetrically aligned information set not always remains asymmetrically aligned after updating it with the best response prices since $\Sigma_1 < S < \Sigma_2 \approx 4.4495$. For the last two values of S an asymmetrically aligned information set always remains asymmetrically aligned as $S > \Sigma_2$.

Concerning the parameters in the learning method, we fix $K = 5$ in the no jump rule. Whenever firms need to pick a price randomly, they use the interval $I = [c, p_H + c]$. We believe that these choices are appropriate since all the prices that are relevant for the long-run outcome of the model lie in interval I and they are always smaller than cK for the model parameters we use, thus the jump size is not restrictive.¹⁶ We consider different values for τ (the number observations used in the estimation). Since there are 4 parameters to be estimated, we need at least 4 observations in the information set. We will investigate how the size of the information set affects the outcome of the simulations.

Since we conjectured to observe substantially different outcomes when firms use all observations compared to the case when they use the last τ observations only, we discuss the simulation results for these cases in separate sections.

4.1 Simulations with all observations

First we investigate the outcome of the model when firms use all observations for estimating the perceived demand function. In this case, firms can move towards a general SSE, they can reach the Nash equilibrium or the ALE (provided it exists). As we have shown, the latter two equilibria are reached only when the initial observations are aligned. Since initial prices are drawn randomly, information sets are typically not aligned, therefore a general SSE is reached, in which firm do not approximate correctly even that linear

¹⁶Also note that our theoretical results do not depend on the rules that augment least squares learning.




Figure 4: Time series of prices for $S = 2$. Other parameters: $s = 1$ and $c = 1$.

	$S = 2$	$S = 4$	$S = 4.2$	$S = 4.35$	$S = 4.5$	$S = 10$
mean	1.6060	1.8648	1.8836	1.9045	1.9269	2.6018
median	1.5781	1.8308	1.8582	1.8787	1.9004	2.5203
stdev	0.1625	0.2372	0.2315	0.2362	0.2451	0.5694

Table 3: Descriptive statistics of final prices for different values of S . Other parameters: $s = 1$ and $c = 1$.

part of the demand function on which they operate.¹⁷

Figure 4 illustrates the time series of prices in a typical simulation for $S = 2$. The figure shows that prices settle down fast and that firms charge different prices. The given simulation stopped in period 1000, the maximal difference between the true and perceived demands at the final price vector is $0.3 \cdot 10^{-3}$, confirming that firms move towards a self-sustaining equilibrium.

As Proposition 3.2 shows, many price vectors can be part of an SSE. Therefore it is worthwhile to investigate the distribution of final prices. Table 3 shows descriptive statistics of the final prices for different values of S . As S increases, both the average and the median prices increase.¹⁸ There is not much difference in the standard deviations. Figure 5 shows histograms of the final prices over the 1000 different initializations, for different values of S . The histograms show that there is substantial price dispersion and that neither the Nash-equilibrium nor the ALE provides a benchmark outcome when all observations are used. As S increases, the distribution seems to become flatter.

¹⁷We need 4×3 initial values for each simulation. We ran numerical simulations to investigate how often initial observations are symmetrically or asymmetrically aligned. Based on 1.000.000 simulations for each value of S we considered, initial observations are symmetrically aligned in less than 0.02% of the cases whereas they are asymmetrically aligned in less than 0.77% of the cases.

¹⁸Note that the upper bound of the interval for initial prices also increases.

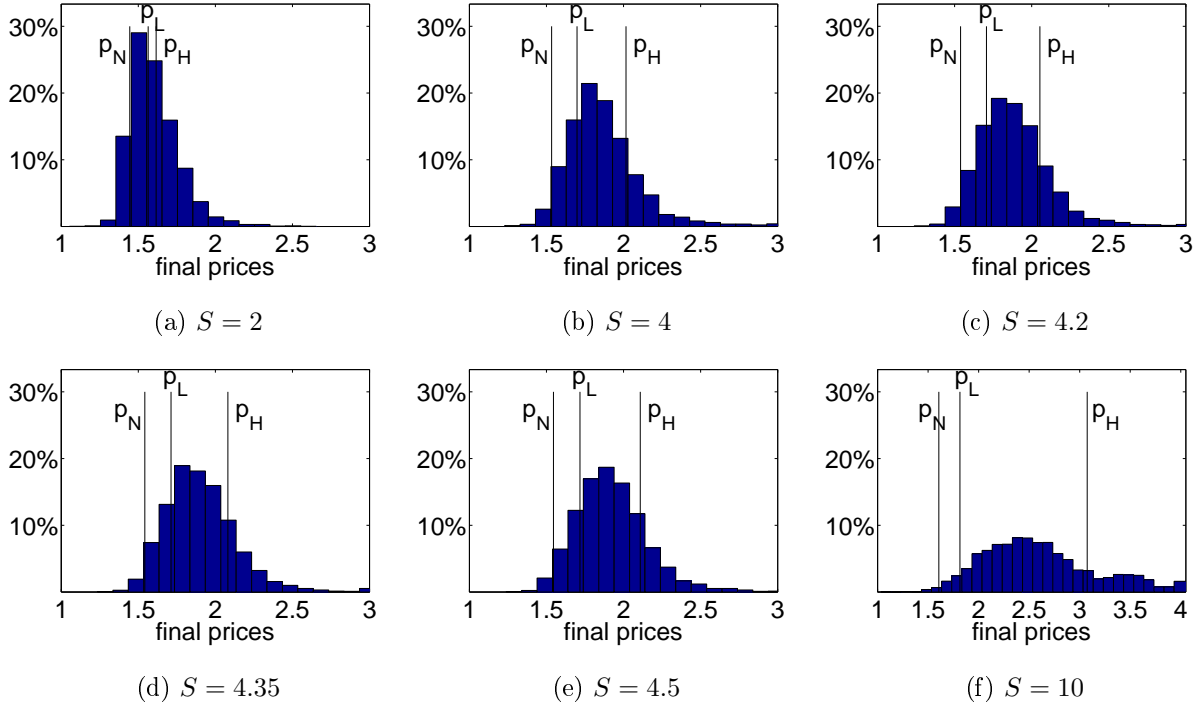


Figure 5: Histogram of final prices for different values of S . Other parameters: $s = 1$ and $c = 1$.

	$S = 2$	$S = 4$	$S = 4.2$	$S = 4.35$	$S = 4.5$	$S = 10$
mean	$5.9 \cdot 10^{-4}$	$8.7 \cdot 10^{-4}$	$8.0 \cdot 10^{-4}$	$8.2 \cdot 10^{-4}$	$7.9 \cdot 10^{-4}$	$1.0 \cdot 10^{-3}$
min	$2.6 \cdot 10^{-8}$	$2.1 \cdot 10^{-16}$	$1.0 \cdot 10^{-15}$	$1.0 \cdot 10^{-15}$	$1.1 \cdot 10^{-16}$	0
max	$4.7 \cdot 10^{-2}$	$1.2 \cdot 10^{-1}$	$8.2 \cdot 10^{-2}$	$9.0 \cdot 10^{-2}$	$7.4 \cdot 10^{-2}$	$9.4 \cdot 10^{-2}$
$\text{diff} \leq 10^{-2}$	978	971	973	974	977	962
$\text{diff} \leq 10^{-3}$	906	895	887	889	887	869
$\text{diff} \leq 10^{-4}$	371	443	454	455	465	615

Table 4: Descriptive statistics of the absolute difference between the true and perceived demands at final prices, for different values of S . Other parameters: $s = 1$ and $c = 1$.

In order to measure how close firms get to a self-sustaining equilibrium, we calculate the absolute difference between the actual and perceived demands at the final price vectors. The difference is 0 in an SSE. Table 4 shows descriptive statistics of these differences for different values of S . The first three rows show the mean, minimal and maximal absolute differences over *individual firms* whereas the last three rows report the number of initializations for which the difference is smaller than 10^{-2} , 10^{-3} and 10^{-4} for the three firms *jointly*.¹⁹

¹⁹For comparison, the mean *initial* difference (i.e. in period 5) ranges from 1 to 2 for the different values of S we consider.

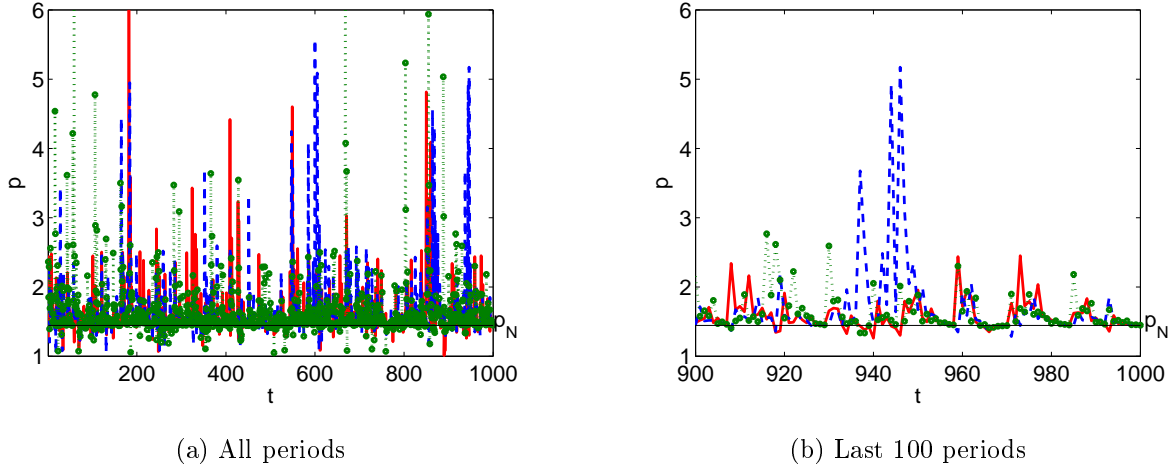


Figure 6: Time series of prices for $\tau = 4$. Parameters: $s = 1$, $S = 2$ and $c = 1$.

We can conclude from the table that differences are rather small in all cases. In almost all cases, the maximal difference is at most 10^{-2} . This confirms that firms get close to a self-sustaining equilibrium when all observations are used in the regression. We practically never observed convergence to the Nash equilibrium or to the ALE.

4.2 Simulations with the last τ observations

Next we turn to the case when information sets contain the last τ observations only. Our conjecture was that information sets become either symmetrically or asymmetrically aligned in this case and firms converge either to the Nash equilibrium or to the asymmetric learning-equilibrium. From Proposition 3.3 we know that the ALE does not exist for $S = 2$ and $S = 4$, thus the Nash equilibrium should always be reached for these values of S .

As we discussed, at least 4 observations are needed for the regression. It turns out that the process does not converge typically when firms use exactly $\tau = 4$ observations. Figure 6 illustrates the time series of prices in a typical simulation with $\tau = 4$. The figure shows that prices do not settle down at the Nash equilibrium price. They are converging towards the Nash equilibrium (already indicating that the Nash equilibrium is locally stable) but every now and then they diverge away from it. The reason behind this is that when there is not enough dispersion in the observations, parameter estimates become imprecise and one of the firms will charge a relatively large price. When firms use 4 observations only, then the weight of a single observation is apparently large enough and the outlier observation can drive the prices far from the equilibrium.

In contrast, when firms use more observations, the weight of a single observation decreases, thus a

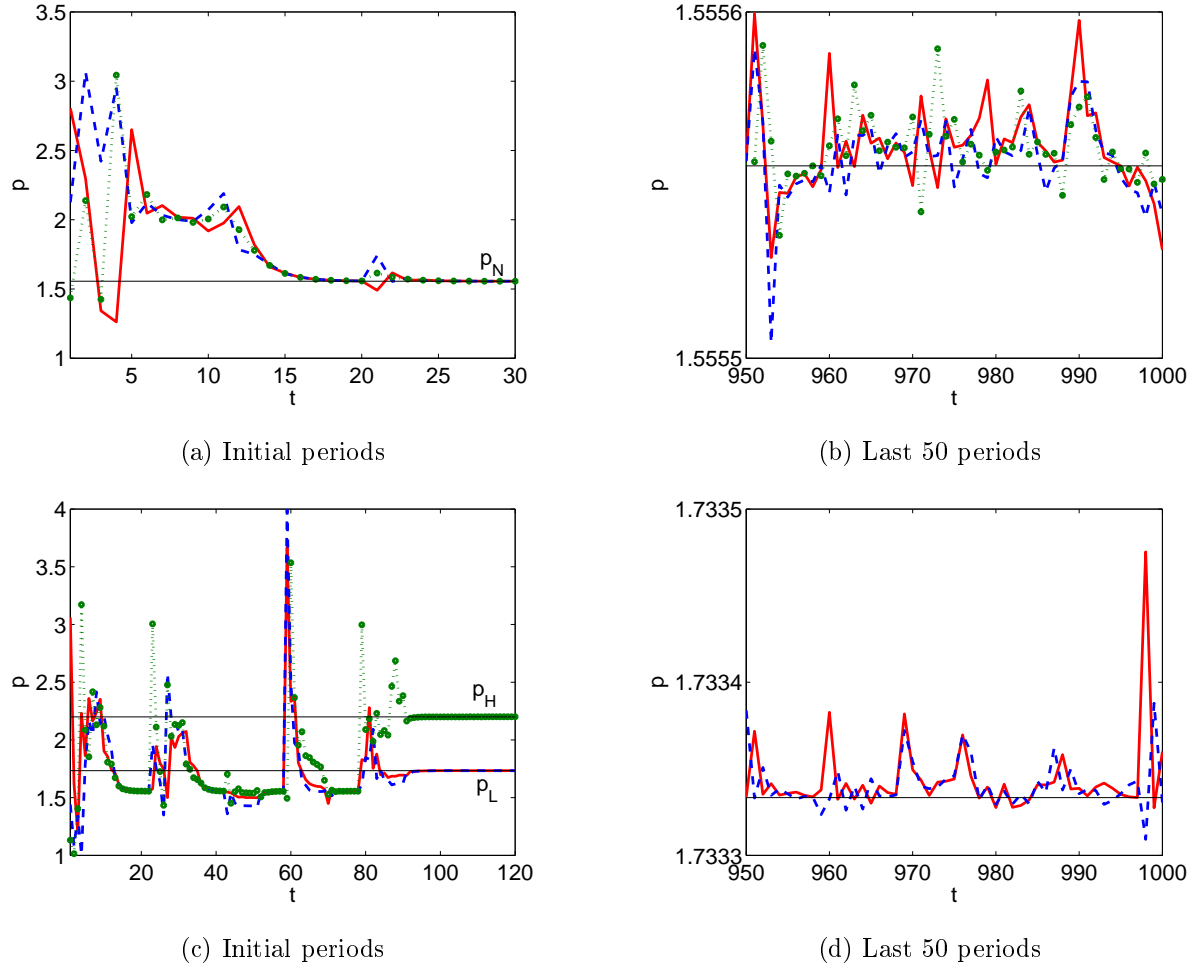


Figure 7: Time series of prices for $\tau = 8$. Parameters: $s = 1$, $S = 5$ and $c = 1$. Panels *a* and *b*: convergence to the Nash equilibrium. Panels *c* and *d*: convergence to the ALE.

single outlier does not drive away prices from the equilibrium that much. We indeed find convergence when the size of the information set increases. Figure 7 shows typical time series for $\tau = 8$. Panel *a* shows an example where prices converge to the Nash equilibrium. Panel *b* shows the same time series but for the last 50 periods of the simulation. It turns out that we do not find exact convergence but small oscillations around the Nash equilibrium. This is caused by the same numerical issue that we have for $\tau = 4$: parameter estimates become imprecise when there is not enough variation in the observations.²⁰ Panels *c* and *d* depict a similar pattern for the case of the ALE.

In order to investigate whether firms always converge either to a neighborhood of the Nash equilibrium

²⁰To confirm that these oscillations are due to numerical issues we run the same simulations with using the true demand coefficients when observations in the information set are aligned. In this case we always find exact convergence to one of the equilibria. These simulations serve as a theoretical benchmark only since the true coefficients are not available for firms.

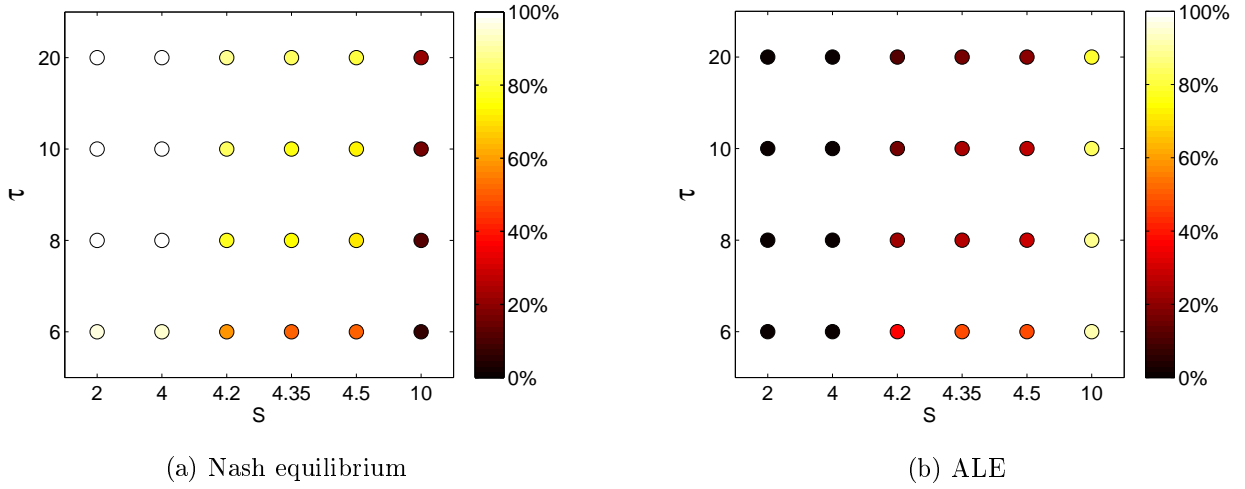


Figure 8: Convergence to the Nash equilibrium (panel a) and to the ALE (panel b) Parameters: $s = 1$ and $c = 1$.

or to a neighborhood of the ALE, we run 1000 simulations for each (S, τ) combination that we consider and we calculate which proportion of the final price vectors lies in a small neighborhood of the Nash equilibrium and the ALE respectively. Figure 8 illustrates the results. For each (S, τ) combination, the figure shows which proportion of the 1000 simulations leads to the 0.001-neighborhood of the Nash equilibrium (panel a) or of the ALE (panel b).²¹ We can conclude from the figure that firms almost always reach either the Nash equilibrium or the ALE, that is prices almost always converge.²² The figure also shows that the Nash equilibrium is reached more often as τ increases. On the other hand, the ALE becomes more dominant as S increases. We will come back to these effects later.

To evaluate the effect of the numerical issue that results in not exact convergence, we ran the same simulations with firms using the true parameters of the given part of the demand function when the information set is aligned. The results (reported in panel a of Table 7 in Appendix B) confirm that firms always reach either the Nash equilibrium or the ALE. We again find that the Nash equilibrium is reached more often as τ increases and that firms converge to the ALE more often as S increases. However, the Nash equilibrium is reached much more often than before. This shows that the numerical issue that occurs when there is not enough variation in the observations, has an important effect on which equilibrium will eventually be reached. The results suggest that the Nash equilibrium is less stable than the ALE in the sense that the numerical issue can drive prices from the Nash equilibrium to the ALE more often than the

²¹We say that a vector (x_1, \dots, x_n) lies in the ε -neighborhood of another vector (y_1, \dots, y_n) if their Euclidean distance is smaller than or equal to ε : $\sqrt{\sum_{i=1}^n (x_i - y_i)^2} \leq \varepsilon$.

²²See Table 5 in Appendix B for the exact values on which the figure is based.

other way around.²³ In fact, panel *c* of Figure 7 shows a situation where prices settle down around the Nash equilibrium initially but after a high price realization the firms converge to the ALE. This finding explains why firms converge more frequently to the Nash equilibrium when τ increases. As we discussed, when the size of the information set increases, a single observation has a smaller effect on the parameter estimates. Therefore a high price realization that may occur after prices have settled down around the Nash equilibrium, has a smaller impact on the parameter estimates, therefore the best-response prices stay in the basin of attraction of the Nash equilibrium instead of reaching the basin of attraction of the ALE.

An alternative way to assess the effect of the numerical issue is to introduce noise in the model. By adding a small noise to the demand faced by firms, there will be enough variation in the observations and therefore parameter estimates do not become imprecise. We drew a random number from the $N(0, 0.001^2)$ distribution in each period for each firm (independently across firms and periods) and added it to the demand faced by the firm. We focused on the case $\tau = 8$ only. The outcomes of these simulations are similar to those when the true coefficients are used in the regression (see panel *b* Table 7 in Appendix B). The reason for this is that the numerical issue that occurs in the model under deterministic demand does not occur under noise and consequently there are no unexpected jumps in prices. This is also confirmed by the fact that there are hardly any switches from the Nash equilibrium to the ALE unlike in the deterministic model.

We checked the robustness of our results with respect to the number of periods in the simulations and the number of different initializations. We focused on the case $\tau = 8$ and we ran two sets of simulations with the previously used values of S : one set with 10000 periods instead of 1000 and another one with 10000 different initializations instead of 1000. The outcomes of these simulations are shown in Table 8 in Appendix B. The outcomes are in line with our previous results. In the case of 10000 runs, we get more or less the same proportions for the different (S, τ) combinations as before.²⁴ Therefore we conclude that our results are robust with respect to the number of simulations. However, the Nash equilibrium is reached less often when simulations ran for 10000 periods.²⁵ This difference is in line with the finding that switches occur from the Nash equilibrium to the ALE but not in the other direction.

²³To investigate whether the Nash equilibrium is indeed less stable than the ALE, we analyzed the entire time series of prices for each initialization and we found that 1. prices leave the 0.001-neighborhood of the Nash more often than that of the ALE and 2. there are barely any switches from the ALE to the Nash equilibrium while there are much more switches in the other direction.

²⁴The mean absolute difference between the proportions is 1.2%-point while the maximal absolute difference is 2.6%-point over the different values of S .

²⁵The mean difference between the proportions for the Nash equilibrium is 3.54%-point while the maximal difference is 8.8%-point over the different values of S .

Finally, we investigated how often the different augmenting rules are used in the simulations. The results show that the random price rule is used a few times but the other two rules are hardly ever used.²⁶ We ran additional simulations with different values of K (the jump size in the no jump rule) as well. The results show that the Nash equilibrium is reached more often as K decreases. The reason for this becomes clear from panel *c* of Figure 7. As we can see, prices settled down around the Nash equilibrium price initially but a large jump around period 60 moved prices towards the ALE-prices. Thus, if we allow for smaller jumps only, prices will be driven out of the neighborhood of the Nash equilibrium less often. But the ALE does not disappear for smaller values of K , we still observe convergence to the ALE as well.

5 Discussion

This paper has focused on learning about market conditions. Firms apply least squares learning where the perceived demand function is correctly specified locally and firms can observe the actions of each other. We have proved that the model has coexisting locally stable equilibria and we have shown that least squares learning can result in a suboptimal outcome for some firms even when the estimated function perfectly fits the market observations. This outcome, the asymmetric learning-equilibrium, was not found in the literature of least squares learning before. Our results show that firms can approximate the demand function better when they use only the most recent observations in the regression. However, the resulting market outcome can be suboptimal for some firms therefore it might worthwhile to experiment with the price (by charging a lower sales price for example) every now and then as this might ensure that firms do not get locked up in the suboptimal situation.

We have considered the Salop model with 3 firms in equidistant locations and with two types of consumers, differing in their transportation costs. Firms do not know the market structure and they apply least squares learning to learn about demand conditions. They approximate the (piecewise linear) demand function with a linear perceived demand function and they maximize their profit subject to their perceived demand function. The model has three kinds of equilibria: a general self-sustaining equilibrium, the Nash equilibrium and the asymmetric learning-equilibrium. In a self-sustaining equilibrium firms approximate the true demand function correctly only in the equilibrium but the approximation is incorrect for out-of-equilibrium prices. In the Nash equilibrium and the ALE firms correctly learn the linear part of the true demand function on which they operate. In the ALE the high-price firm underestimates the demand for low prices and it attracts high-type consumers only. This situation is suboptimal for the firm when the two transportation costs are not too different. However, when transportation costs become sufficiently different,

²⁶See Table 6 in Appendix B for more details.

then the high-price firm is better-off by not attracting low-type consumers. The ALE has worse welfare properties than the Nash equilibrium outcome. We have proved that both the Nash equilibrium and the ALE are locally stable, thus the model can have coexisting locally stable equilibria. When firms use all past observations in the approximation, then they typically reach a general SSE. On the other hand, when only the most recent observations are used, firms converge either towards the Nash equilibrium or towards the ALE. As firms use more observations in the regression (but not all observations), the Nash equilibrium is reached more often. In contrast, the ALE is reached more often as the transportation cost of high-type consumers increases.

In the model we have made some assumptions whose effects should be discussed. First of all, we have introduced heterogeneity on the demand side of the market. This is not an unrealistic assumption as consumers could easily differ in their transportation costs, moreover it makes the model more general. With homogeneous consumers, the true demand function is still piecewise linear so least squares learning can lead to an SSE or to the Nash equilibrium. The ALE, however, does not exist since if a firm does not attract any consumers, then it will charge a lower price eventually since the observations with zero demand will move the parameter estimates in a direction that yields a lower price. Thus, consumer heterogeneity is essential for having an asymmetric outcome.

We have assumed that firms do not know anything about the demand structure. This assumption might sound too restrictive in the transportation cost interpretation of the Salop model. On the other hand, if we consider the product space interpretation of the circular road, then it is not unreasonable to assume that firms do not know how consumers evaluate different product types. From a different point of view, the method that the firms are using to learn about demand conditions is extremely natural. They do not assume a complicated structural model behind the data, instead they are using the simplest possible functional form in the regression. When they do not use all observations in the regression, they indeed find that the linearity assumption is correct so there is no need to reconsider the functional form. Nevertheless, the assumption of not knowing the demand structure is crucial for the results: If firms knew the underlying demand structure, they could recover the transportation costs and the amount of each consumer type from the parameter estimates and they could find the Nash equilibrium of the static game.

The assumption of symmetry between firms (equidistant locations, equal marginal costs of production) can be relaxed to some extent. If locations are not far from each other and marginal costs are not very different, then the model has the same types of equilibria. The ALE becomes even more plausible in the sense that all three firms charge different prices thus it will not be suspicious for the high-price firm that the other two firms charge the same price. We assumed symmetry only to simplify calculations.

Let us now elaborate on whether our results could still hold in more general models. If we increase

the number of firms or the number of consumer types in the Salop framework, the true demand function remains piecewise linear and thus firms may still learn at most one linear part correctly. Therefore firms can reach an SSE or the Nash equilibrium again. Preliminary simulations confirm that asymmetric outcomes can occur for higher numbers of firms too. In fact, there are multiple asymmetric outcomes: These outcomes differ with respect to the number as well as the locations of high-price firms. We expect that the same kind of outcomes can occur in different market structures too. Consider for example a Bertrand oligopoly with heterogeneous goods. Suppose that consumers are heterogeneous and that two different functions characterize their demand. This situation can lead to a kink in the aggregate demand functions. If the functional form of the perceived demand function coincides with the functional form of one (or both) part of the true demand functions, then the same kind of asymmetric equilibrium may easily occur. Some firms might focus on one consumer group only while other firms might attract both consumer groups.

References

- Anufriev, M., Kopányi, D., and Tuinstra, J. (2013). Learning cycles in Bertrand competition with differentiated commodities and competing learning rules. *Journal of Economic Dynamics and Control*, 37(12):2562–2581.
- Blume, L. E., Bray, M. M., and Easley, D. (1982). Introduction to the stability of rational expectations equilibrium. *Journal of Economic Theory*, 26(2):313–317.
- Blume, L. E. and Easley, D. (1982). Learning to be rational. *Journal of Economic Theory*, 26(2):340–351.
- Bonanno, G. (1988). Oligopoly equilibria when firms have local knowledge of demand. *International Economic Review*, 29(1):45–55.
- Bonanno, G. and Zeeman, E. C. (1985). Limited knowledge of demand and oligopoly equilibria. *Journal of Economic Theory*, 35(2):276–283.
- Bray, M. (1982). Learning, estimation, and the stability of rational expectations. *Journal of Economic Theory*, 26(2):318–339.
- Brousseau, V. and Kirman, A. (1992). Apparent convergence of learning processes in mis-specified games. In Dutta, B., (editor), *Game Theory and Economic Applications*. Springer-Verlag.
- Evans, G. and Honkapohja, S. (2001). *Learning and Expectations in Macroeconomics*. Princeton University Press.

- Gates, D., Rickard, J., and Wilson, D. (1977). A convergent adjustment process for firms in competition. *Econometrica*, 45(6):1349–1363.
- Kirman, A. (1983). On mistaken beliefs and resultant equilibria. In Frydman, R. and Phelps, E. (editors), *Individual Forecasting and Collective Outcomes*. Cambridge University Press.
- Kirman, A. P. (1975). Learning by firms about demand conditions. In Day, R. H. and Groves, T. (editors), *Adaptive Economic Models*, pages 137–156. Academic Press, New York.
- Kopányi, D. (2013). Heterogeneous learning in Bertrand competition with differentiated goods. In Teglio, A., Alfarano, S., Camacho-Cuena, E., and Ginés-Vilar, M. (editors), *Managing market complexity: The approach of artificial economics*. Springer-Verlag.
- Marcet, A. and Sargent, T. (1989). Convergence of least squares learning mechanisms in self-referential linear stochastic models. *Journal of Economic Theory*, 48(2):337–368.
- Salop, S. C. (1979). Monopolistic competition with outside goods. *The Bell Journal of Economics*, 10(1):141–156.
- Schinkel, M. P., Tuinstra, J., and Vermeulen, D. (2002). Convergence of Bayesian learning to general equilibrium in mis-specified models. *Journal of Mathematical Economics*, 38(4):483–508.
- Silvestre, J. (1977). A model of general equilibrium with monopolistic behavior. *Journal of Economic Theory*, 16(2):425–442.
- Tirole, J. (1988). *The theory of industrial organization*. MIT press.
- Tuinstra, J. (2004). A price adjustment process in a model of monopolistic competition. *International Game Theory Review*, 6(3):417–442.
- Weddepohl, C. (1995). A cautious price adjustment mechanism: chaotic behavior. *Journal of Economic Behavior & Organization*, 27(2):293–300.

APPENDIX

A Proofs and derivations

A.1 The proof of Proposition 2.1

PROOF.

First note that $p_i \geq c$ must hold for each firm in equilibrium. Otherwise the firm with the lowest price, say firm j , would always face a positive demand and would make a certain loss on each product. The firm could increase its profit by choosing a higher price for which its profit is at least 0. This can be achieved by any $p_j \geq c$.

Now we will show that each firm must face a positive demand in equilibrium. To see this suppose that firm i is driven out of the whole market by firm j , that is $p_j < p_i - \frac{1}{3}S$. Let $c \leq p_j \leq p_k$ without loss of generality. In this case, firm i can increase its profit by choosing the price $p_i = c + \varepsilon$ for a sufficiently small but positive ε . For this price firm i will not be driven out of the market since $p_i - \frac{1}{3}S = c + \varepsilon - \frac{1}{3}S < c$ for a sufficiently small ε , meaning that firm i can only be driven out of the market with a price that is smaller than the marginal cost. This, as we have seen, cannot occur in equilibrium.

The condition that each firm must have a positive demand in equilibrium implies that all three firms must attract high-type consumers. Thus, equilibria can differ only in the number of firms attracting low-type consumers. There might be three possibilities: 3, 2 or 1 firm attracts low-type consumers. We investigate these cases separately.

Case 1: symmetric Nash equilibrium

When all three firms attract low-type consumers, then firm i faces the following demand function: $D_i(p) = \frac{2}{3} + (p_j + p_k - 2p_i) \left(\frac{1}{2S} + \frac{1}{2s} \right)$. To see this note that there is one low-type and one high-type indifferent consumer between any two firms. The low-type indifferent consumer between firms i and j is at the distance $x = \frac{p_j - p_i}{2s} + \frac{1}{6}$ from firm i . A similar formula applies for the high-type indifferent consumer and for the indifferent consumers between firms i and k .

Firm i maximizes its profit with respect to its price: $\max_{p_i} (p_i - c)D_i(p)$. The first-order conditions²⁷

²⁷Note that the profit functions are quadratic, with the quadratic terms having a negative coefficient. Thus, the second-order conditions are satisfied.

for firms 1, 2 and 3 respectively are

$$\frac{2}{3} + (p_2 + p_3 - 2p_1) \left(\frac{1}{2S} + \frac{1}{2s} \right) - \left(\frac{1}{S} + \frac{1}{s} \right) (p_1 - c) = 0, \quad (\text{A.1})$$

$$\frac{2}{3} + (p_1 + p_3 - 2p_2) \left(\frac{1}{2S} + \frac{1}{2s} \right) - \left(\frac{1}{S} + \frac{1}{s} \right) (p_2 - c) = 0, \quad (\text{A.2})$$

$$\frac{2}{3} + (p_1 + p_2 - 2p_3) \left(\frac{1}{2S} + \frac{1}{2s} \right) - \left(\frac{1}{S} + \frac{1}{s} \right) (p_3 - c) = 0. \quad (\text{A.3})$$

Subtracting (A.2) from (A.1) yield $\frac{5}{2}(p_2 - p_1) \left(\frac{1}{S} + \frac{1}{s} \right) = 0$, from which $p_1 = p_2$. Similarly, subtracting (A.3) from (A.2) gives $p_1 = p_3$. Let p_N denote this common price. Then the first-order conditions simplify to $\frac{2}{3} - \left(\frac{1}{S} + \frac{1}{s} \right) (p_N - c) = 0$, from which

$$p_N = \frac{2Ss}{3(S+s)} + c.$$

The corresponding profits are $\pi_N = \frac{2}{3}(p_N - c) = \frac{4}{9} \frac{Ss}{S+s}$.

The price vector $p = (p_N, p_N, p_N)$ constitutes a Nash equilibrium only if none of the firms has an incentive to deviate from this price unilaterally. A firm can deviate in two possible ways. It can drive out the two other firms from the market of the low-type consumers or it can drive out the other firms from the whole market.²⁸

Let us first consider the case when firm 1 chooses $p_N - \frac{1}{3}S \leq p_1 \leq p_N - \frac{1}{3}s$. In this case firm 1 attracts the low-type consumers and the three firms share the high-type consumers. Thus, firm 1 faces the following demand function: $D_1(p) = \frac{4}{3} + \frac{p_N - p_1}{S}$. To find the optimal price, the following constrained optimization problem needs to be solved:

$$\max_{p_1 \leq p_N - \frac{s}{3}} (p_1 - c) \left(\frac{4}{3} + \frac{p_N - p_1}{S} \right).$$

The Karush-Kuhn-Tucker conditions yield

$$\begin{aligned} \frac{4}{3} + \frac{p_N - p_1}{S} - \frac{1}{S}(p_1 - c) &\geq 0 \\ \left(\frac{4}{3} + \frac{p_N - p_1}{S} - \frac{1}{S}(p_1 - c) \right) \left(p_N - \frac{s}{3} - p_1 \right) &= 0. \end{aligned}$$

Let us suppose that $\frac{4}{3} + \frac{p_N - p_1}{S} - \frac{1}{S}(p_1 - c) = 0$. This gives $p_1 = \frac{2S}{3} + \frac{p_N + c}{2}$. We need to check whether the

²⁸We do not have to consider marginal deviations from p_N since the first-order conditions imply that the local profit maximum is reached at p_N .

condition $p_1 \leq p_N - \frac{s}{3}$ is satisfied.

$$\begin{aligned}
\frac{2S}{3} + \frac{p_N + c}{2} &\leq p_N - \frac{s}{3} \\
0 &\leq \frac{p_N - c}{2} - \frac{2S}{3} - \frac{s}{3} \\
0 &\leq \frac{Ss}{3(S+s)} - \frac{2S}{3} - \frac{s}{3} \\
0 &\leq Ss - 2S(S+s) - s(S+s) \\
0 &\leq -2S^2 - 2Ss - s^2,
\end{aligned}$$

where we used the formula for p_N . The last condition is never satisfied so we can conclude that $p_1^D = p_N - \frac{1}{3}s$ is the optimal deviation in this case. The corresponding demand and profit are $q_1^D = \frac{1}{3}(4 + \frac{s}{S})$ and $\pi_1^D = (p_N - \frac{1}{3}s - c) \frac{1}{3}(4 + \frac{s}{S})$, which simplifies to $\pi_1^D = \frac{s}{9} \frac{S-s}{S+s} (4 + \frac{s}{S})$. Firm 1 does not have an incentive to deviate if $\pi_N \geq \pi_1^D$, which gives

$$\begin{aligned}
\frac{4}{9} \frac{Ss}{S+s} &\geq \frac{s}{9} \frac{S-s}{S+s} \left(4 + \frac{s}{S}\right) \\
4S &\geq (S-s) \left(4 + \frac{s}{S}\right) \\
0 &\geq -3s - \frac{s^2}{S}.
\end{aligned}$$

The last inequality is always satisfied as $S, s > 0$. Thus, this deviation is never profitable.

Now let us consider the other deviation when firm 1 drives the other firms out of the whole market. In this case $p_1 < p_N - \frac{1}{3}S$ should hold. Note, however, that $p_N - \frac{1}{3}S = \frac{2Ss}{3(S+s)} + c - \frac{1}{3}S = \frac{1}{3}S \frac{s-S}{S+s} + c < c$ as $s - S < 0$. This means that firm 1 would have to charge a price below the marginal cost to attract every consumer, leading to negative profits.

Thus, firms do not have an incentive to deviate unilaterally from the price p_N . The price vector $p = (p_N, p_N, p_N)$ is the unique symmetric Nash equilibrium.

Case 2: asymmetric situation with 2 firms serving low-type consumers

Now we will show that the situation in which exactly one firm focuses only on high-type consumers, cannot constitute a Nash equilibrium. Assume without loss of generality that firm 3 charges a high price such that only high-type consumers buy from firm 3: $\min\{p_1, p_2\} + \frac{S}{3} \geq p_3 \geq \min\{p_1, p_2\} + \frac{s}{3}$. In this situation

the demand functions are as follows:

$$D_1(p) = \frac{5}{6} + \frac{p_2 - p_1}{s} + \frac{p_2 + p_3 - 2p_1}{2S}, \quad (\text{A.4})$$

$$D_2(p) = \frac{5}{6} + \frac{p_1 - p_2}{s} + \frac{p_1 + p_3 - 2p_2}{2S}, \quad (\text{A.5})$$

$$D_3(p) = \frac{1}{3} + \frac{p_1 + p_2 - 2p_3}{2S}. \quad (\text{A.6})$$

Profit maximization yields the following first-order conditions (for firms 1, 2 and 3 respectively)

$$\frac{5}{6} + \frac{p_2 - p_1}{s} + \frac{p_2 + p_3 - 2p_1}{2S} - \left(\frac{1}{s} + \frac{1}{S} \right) (p_1 - c) = 0, \quad (\text{A.7})$$

$$\frac{5}{6} + \frac{p_1 - p_2}{s} + \frac{p_1 + p_3 - 2p_2}{2S} - \left(\frac{1}{s} + \frac{1}{S} \right) (p_2 - c) = 0, \quad (\text{A.8})$$

$$\frac{1}{3} + \frac{p_1 + p_2 - 2p_3}{2S} - \frac{1}{S} (p_3 - c) = 0. \quad (\text{A.9})$$

If a Nash equilibrium exists in the given situation, it must be the solution of these first-order conditions. By subtracting (A.8) from (A.7), it can be seen that $p_1 = p_2$ must hold. Therefore, let $p_1 = p_2 = p_L$ and $p_3 = p_H$. The first-order conditions then simplify to

$$\frac{5}{6} + \frac{p_H - p_L}{2S} - \left(\frac{1}{s} + \frac{1}{S} \right) (p_L - c) = 0, \quad (\text{A.10})$$

$$\frac{1}{3} + \frac{p_L - p_H}{S} - \frac{1}{S} (p_H - c) = 0. \quad (\text{A.11})$$

Subtracting (A.11) from (A.10) yields $\frac{1}{2} + \frac{3}{2} \frac{p_H - p_L}{S} - \frac{1}{s} (p_L - c) + \frac{1}{S} (p_H - p_L) = 0$, from which

$$\frac{p_H - p_L}{2S} = \frac{p_L - c}{5s} - \frac{1}{10}. \quad (\text{A.12})$$

Combining (A.10) with (A.12) gives $\frac{11}{15} + \frac{p_L - c}{5s} - \left(\frac{1}{s} + \frac{1}{S} \right) (p_L - c) = 0$. Solving this equation for p_L yields

$$p_L = \frac{11Ss}{12S + 15s} + c.$$

Plugging this expression for p_L in (A.12) yields an equation that can be solved for p_H . The solution simplifies to

$$p_H = \frac{2S^2 + 8Ss}{12S + 15s} + c.$$

Note that the previous calculations yield admissible prices only when $p_L + \frac{1}{3}S \geq p_H \geq p_L + \frac{1}{3}s$, or equivalently $\frac{1}{3}s \leq p_H - p_L \leq \frac{1}{3}S$. Using that

$$p_H - p_L = \frac{2S^2 - 3Ss}{12S + 15s},$$

the condition $\frac{1}{3}s \leq p_H - p_L$ leads to $4Ss + 5s^2 \leq 2S^2 - 3Ss$, or equivalently $2 \left(\frac{S}{s} \right)^2 - 7 \frac{S}{s} - 5 \geq 0$. Solving this quadratic equation gives that $\frac{S}{s} \geq \frac{7 + \sqrt{89}}{4}$ must hold.

The condition $p_H - p_L \leq \frac{1}{3}S$ leads to $12S + 15s \geq 6S - 9s$, from which $6S + 24s \geq 0$. This condition is satisfied as $S, s > 0$. Thus, this type of asymmetric Nash equilibrium may exist only when $\frac{S}{s} \geq \frac{7+\sqrt{89}}{4}$.

The price vector $p = (p_L, p_L, p_H)$ constitutes a Nash equilibrium only if none of the firms has an incentive to deviate unilaterally. Now we will show that either the high-price firm or the low-price firms can earn a higher profit by charging a different price. First, let us calculate the profits under $p = (p_L, p_L, p_H)$. Plugging the prices in demand functions (A.4)-(A.6) yields $q_1 = q_2 = q_L = \frac{11S+11s}{12S+15s}$ and $q_3 = q_H = \frac{2S+8s}{12S+15s}$. The corresponding profits are $\pi_1 = \pi_2 = \pi_L = \frac{121Ss(S+s)}{(12S+15s)^2}$ and $\pi_3 = \pi_H = S \left(\frac{2S+8s}{12S+15s} \right)^2$.

First let us suppose that the high-price firm deviates and charges $p_3^D = p_L$, where superscript D refers to deviation. In that case $q_3^D = \frac{2}{3}$ since all three firms charge the same price. The corresponding profit is $\pi_3^D = \frac{2}{3} \frac{11Ss}{12S+15s}$. This deviation leads to a higher profit for firm 3 when

$$\begin{aligned} \frac{2}{3} \frac{11Ss}{12S+15s} &> S \left(\frac{2S+8s}{12S+15s} \right)^2 \\ 11s(12S+15s) &> 6S^2 + 48Ss + 96s^2 \\ 0 &> 6S^2 - 84Ss - 69s^2. \end{aligned} \quad (\text{A.13})$$

Now let us suppose that firm 1 deviates by charging $p_1^D = p_H$. In that case firm 1 serves the high-type consumers only so it faces a similar demand function as (A.6). Thus, its demand equals $q_1^D = \frac{1}{3} + \frac{p_L - p_H}{2S} = \frac{6S+13s}{2(12S+15s)}$ and the corresponding profit is $\pi_1^D = \frac{6S+13s}{2(12S+15s)} \frac{S(S+4s)}{12S+15s}$. This deviation leads to a higher profit for firm 1 when

$$\begin{aligned} \frac{6S+13s}{12S+15s} \frac{S(S+4s)}{12S+15s} &> \frac{121Ss(S+s)}{(12S+15s)^2} \\ (6S+13s)(S+4s) &> 121s(S+s) \\ 6S^2 - 84Ss - 69s^2 &> 0. \end{aligned} \quad (\text{A.14})$$

Comparing conditions (A.13) and (A.14), we find that one of the firms always has an incentive to deviate whenever $6S^2 - 84Ss - 69s^2 \neq 0$. Now we will show that the high-price firm has an incentive to deviate even if the previous equation holds with equality. Note that we did not consider the *optimal* deviation in the previous calculations. We only showed that there exists a deviation that is more profitable under certain conditions. When $6S^2 - 84Ss - 69s^2 = 0$ holds, firm 3 is indifferent between charging p_L and p_H (keeping the price of the other two firms fixed):

$$\pi_3(p_L, p_L, p_L) = \pi_3(p_L, p_L, p_H). \quad (\text{A.15})$$

We will now show that the marginal profit of firm 3 is not equal to 0 at $p = (p_L, p_L, p_L)$. This implies that a marginal deviation from $p_3 = p_L$ (in the appropriate direction) yields a strictly higher profit, thus $p = (p_L, p_L, p_H)$ cannot be a Nash equilibrium.

The marginal profit of firm 3 at $p = (p_L, p_L, p_L)$ can be calculated using (A.3):

$$\left. \frac{\partial \pi_3}{\partial p_3} \right|_{p=(p_L, p_L, p_L)} = \frac{2}{3} - \left(\frac{1}{S} + \frac{1}{s} \right) (p_L - c).$$

Plugging in the formula for p_L yields $\frac{2}{3} - \frac{S+s}{Ss} \frac{11Ss}{12S+15s}$, which simplifies to $\frac{7S+9s}{12S+15s}$. This expression is always positive since $S, s > 0$. Thus, firm 3 can get a strictly higher profit by marginally increasing its price: $\pi_3(p_L, p_L, p_L + \varepsilon) > \pi_3(p_L, p_L, p_L)$ for a small enough $\varepsilon > 0$. Combining the last inequality with (A.15) shows that $p = (p_L, p_L, p_H)$ cannot be a Nash equilibrium.

Thus, we have shown that one of the firms can always get a higher profit by unilaterally changing its price. We can conclude that there does not exist an asymmetric Nash equilibrium in pure strategies where exactly two firms attract low-type consumers.

Case 3: asymmetric situation with 1 firm serving low-type consumers

Now we will show that the situation in which two firms focus only on the high-type consumers, cannot constitute a Nash equilibrium. Assume without loss of generality that firm 1 charges a low price such that it attracts every low-type consumer: $p_1 + \frac{s}{3} \geq \{p_2, p_3\} \geq p_1 + \frac{s}{3}$. In this situation the demand functions are as follows:

$$\begin{aligned} D_1(p) &= \frac{4}{3} + \frac{p_2 + p_3 - 2p_1}{2S}, \\ D_2(p) &= \frac{1}{3} + \frac{p_1 + p_3 - 2p_2}{2S}, \\ D_3(p) &= \frac{1}{3} + \frac{p_1 + p_2 - 2p_3}{2S}, \end{aligned}$$

with the corresponding first-order conditions for profit maximization

$$\frac{4}{3} + \frac{p_2 + p_3 - 2p_1}{2S} - \frac{1}{S}(p_1 - c) = 0, \quad (\text{A.16})$$

$$\frac{1}{3} + \frac{p_1 + p_3 - 2p_2}{2S} - \frac{1}{S}(p_2 - c) = 0, \quad (\text{A.17})$$

$$\frac{1}{3} + \frac{p_1 + p_2 - 2p_3}{2S} - \frac{1}{S}(p_3 - c) = 0. \quad (\text{A.18})$$

By subtracting (A.18) from (A.17), it can be seen that $p_2 = p_3$ must hold. Let $p_1 = p_L$ and $p_2 = p_3 = p_H$. Then the first-order conditions simplify to

$$\frac{4}{3} + \frac{p_H - p_L}{S} - \frac{1}{S}(p_L - c) = 0, \quad (\text{A.19})$$

$$\frac{1}{3} + \frac{p_L - p_H}{2S} - \frac{1}{S}(p_H - c) = 0. \quad (\text{A.20})$$

Subtracting (A.20) from (A.19) yields $1 + \frac{3}{2S}(p_H - p_L) + \frac{1}{S}(p_H - p_L) = 0$. This equation, however, does not give an admissible solution. Since every coefficient is positive and the right hand side is 0, $p_H < p_L$ must hold, which contradicts the assumption $p_H \geq p_L + \frac{s}{3}$. Thus, there exists no asymmetric pure-strategy Nash equilibrium in which exactly one firm serves the low-type consumers. \square

A.2 The proof of Proposition 3.2

PROOF.

To simplify notation, let $D_i^P(p) = A_i - b_{ii}p_i$, where $A_i = a_i + b_{ij}p_j + b_{ik}p_k$. Then using (3) the best response price is given by

$$p_i^{\text{BR}} = \frac{A_i}{2b_{ii}} + \frac{c}{2}. \quad (\text{A.21})$$

Since $p_i = p_i^{\text{BR}}$ in an SSE, the perceived demand is given by

$$D_i^P(p) = \frac{A_i - b_{ii}c}{2}. \quad (\text{A.22})$$

Note that 9 variables characterize an SSE under the simplified notation: 1 price and the 2 parameters of the perceived demand function for each firm. On the other hand, there are 6 conditions (best response price and equality of actual and perceived demands for each firm). Thus, the system of equations that characterizes an SSE might be solved, with 3 free variables. We will now show that for a given price vector $p = (p_i, p_j, p_k)$ we can find values of $\{A_i, b_{ii}\}_{i=1}^3$ such that the system is in an SSE.

From (A.21) we get $A_i = b_{ii}(2p_i - c)$. Combining this with (A.22), the perceived demand simplifies to $D_i^P(p) = b_{ii}(p_i - c)$. Since the actual and the perceived demands must coincide at price vector $p = (p_i, p_j, p_k)$, it must hold that $D_i(p) = b_{ii}(p_i - c)$, from which

$$b_{ii} = \frac{D_i(p)}{p_i - c}. \quad (\text{A.23})$$

Combining this with the previous formula for A_i yields

$$A_i = \frac{D_i(p)}{p_i - c} (2p_i - c). \quad (\text{A.24})$$

Thus, for a given price vector $p = (p_i, p_j, p_k)$, formulas (A.23) and (A.24) specify the values of b_{ii} and A_i under which the system is in an SSE.

Let us investigate which price vectors lead to an economically sensible perceived demand function. That is, we want to characterize the set of prices for which $b_{ii} > 0$ and $A_i > 0$ (i.e. the perceived demand function is downward-sloping and the “intercept” is positive)²⁹.

It follows from (A.23) that $b_{ii} > 0$ if and only if $D_i(p) > 0$ and $p_i > c$. Under these conditions, $A_i > 0$ is satisfied as well. \square

²⁹For having an economically sensible perceived demand function, one might consider introducing the conditions $a_i > 0$,

A.3 The proof of Proposition 3.3

PROOF.

We know from Proposition 3.2 that $D_i(p) > 0$ must hold for each firm. Since each firm must face a positive demand in an SSE, each firm must attract high-type consumers. This implies that there are three possible SSE in which firms correctly learn one linear part of the true demand function, depending on whether 1, 2 or 3 firms serve low-type consumers.³⁰

When all 3 firms attract low-type consumers, then demand conditions are characterized by $D_i(p) = \frac{2}{3} + (p_j + p_k - 2p_i) \left(\frac{1}{2S} + \frac{1}{2s} \right)$ (see Case 1 in the proof of Proposition 2.1). The best response function can be derived from first-order conditions (A.1)-(A.3). As we have seen, these first-order conditions have a unique solution, which corresponds to the Nash equilibrium of the model with known demand. Thus, when all 3 firms serve both consumer types and firms correctly learn the corresponding linear part of the true demand function, then the Nash equilibrium is the unique steady state of the learning process.

When only 2 firms attract low-type consumers, then demand conditions are characterized by (A.4)-(A.6) (see Case 2 in the proof of Proposition 2.1). The corresponding best response functions can be derived from first-order conditions (A.7)-(A.9). These first-order conditions have a unique solution, in which the low-price firms charge $p_L = \frac{11Ss}{12S+15s} + c$ and the high-price firm asks the price $p_H = \frac{2S^2+8Ss}{12S+15s}$. We have also seen that this outcome exists only when $\frac{S}{s} \geq \Sigma_1$. Even though this outcome is not a Nash equilibrium of the model with known demand, it is a steady state of the learning process. The reason behind this is that firms do not know that it would be profitable to change their price unilaterally since they approximate the demand function with a linear function, implying that they do not know that they would get a much higher demand by undercutting other firms. Thus, when only 2 firms serve both consumer types and firms correctly learn the corresponding linear part of the true demand function, then the unique steady state is given by 2 firms charging p_L and 1 firm charging p_H . We refer to this outcome as asymmetric learning-equilibrium.

We have seen that when only 1 firm serves the low-type consumer, then first-order conditions (A.16)-(A.18) do not yield an admissible solution. Therefore the learning process does not have a steady state in this situation.

This shows that the Nash equilibrium and the asymmetric learning-equilibrium are the only steady states in which all three firms correctly learn the linear part of the true demand function on which they $b_{ij} > 0$ and $b_{ik} > 0$ in addition to the condition $A_i > 0$. Note, however, that these extra conditions do not restrict the set of admissible prices further as for a given positive A_i one can always find values for a_i , b_{ij} and b_{ik} such that $A_i = a_i + b_{ij}p_j + b_{ik}p_k$ holds and the conditions on the signs are satisfied.

³⁰Note that these are exactly the same cases that we analyzed in the proof of Proposition 2.1.

operate. □

A.4 The proof of Proposition 3.4

First we show that the Nash equilibrium price is smaller than the lower price in the ALE. Comparing p_N and p_L , we get that $p_N < p_L$ if and only if $\frac{2Ss}{3(S+s)} < \frac{11Ss}{12S+15s}$. This reduces to $0 < 9S + 3s$, which always holds since $S, s > 0$. We have shown before that $p_H > p_L$ whenever the ALE exists. Thus, we have $p_N < p_L < p_H$.

We have seen in the proof of Proposition 2.1 that the Nash-equilibrium profit is $\pi_N = \frac{4}{9} \frac{Ss}{S+s}$ while the profits in the ALE are given by $\pi_L = \frac{121Ss(S+s)}{(12S+15s)^2}$ and $\pi_H = S \left(\frac{2S+8s}{12S+15s} \right)^2$. The low-price firms make a higher profit than the high-price firm only if $121s(S+s) > 4(S+4s)^2$, from which $0 > 4 \left(\frac{S}{s} \right)^2 - 89 \frac{S}{s} - 57$. This gives $\frac{S}{s} < \frac{89+11\sqrt{73}}{8} \approx 22.87$.

The Nash-equilibrium profit is always smaller than the profit of low-price firms in the ALE: $\frac{4}{9} \frac{Ss}{S+s} < \frac{121Ss(S+s)}{(12S+15s)^2}$ if and only if $4(12S+15s)^2 < 1089(S+s)^2$, which reduces to $0 < 9S + 3s$. This inequality is always satisfied.

Next we show that the Nash-equilibrium profit is larger than the profit of the high-price firm in the ALE only if $\frac{S}{s}$ is low enough. $\frac{4}{9} \frac{Ss}{S+s} > S \left(\frac{2S+8s}{12S+15s} \right)^2$ if and only if $(4S+5s)^2 s > (S+s)(S+4s)^2$. This is equivalent to the following inequality: $-\left(\frac{S}{s}\right)^3 + 7\left(\frac{S}{s}\right)^2 + 16\frac{S}{s} + 9 > 0$. First we will argue that the function $f(x) = -x^3 + 7x^2 + 16x + 9$ has a single real root. Note that f is a cubic function, therefore it may have 1, 2 or 3 real roots. It is easy to see that f has a local maximum at $x_+ = \frac{7+\sqrt{97}}{3} \approx 5.61$ and a local minimum at $x_- = \frac{7-\sqrt{97}}{3} \approx -0.95$. Since the function value is positive both at the local maximum and at the local minimum (around 142.51 and 0.97, respectively), the function has a unique root. Numerical calculations show that this root is around 8.91. Thus, when $\frac{S}{s} < 8.91$, then $\pi_N > \pi_H$, otherwise the opposite relation holds.

Finally, we show that the total profit of firms in the Nash equilibrium is always smaller than in the ALE. Using the previous formulas, $3\pi_N < 2\pi_L + \pi_H$ reduces to $6s(4S+5s)^2 < 121s(S+s)^2 + 2(S+4s)^2(S+s)$. This simplifies further to $0 < 2S^3 + 43S^2s + 50Ss^2 + 3s^3$. This inequality is always satisfied as $S, s > 0$.

A.5 The proof of Lemma 3.6

PROOF.

When price observations are aligned, estimation yields the true parameters that characterize the given linear part of the demand function. Under symmetrically aligned price observations the parameter estimates are given by $a_i = \frac{2}{3}$, $b_{ii} = \frac{1}{S} + \frac{1}{s}$ and $b_{ij} = b_{ik} = \frac{1}{2S} + \frac{1}{2s}$ (see Case 1 in the proof of Proposition 2.1 for

the corresponding demand function). Thus, using (3), the best response of firm i is

$$p_i^{\text{BR}} = \frac{\frac{2}{3} + \left(\frac{1}{2S} + \frac{1}{2s}\right)(p_j + p_k)}{2\left(\frac{1}{S} + \frac{1}{s}\right)} + \frac{c}{2}.$$

Then $|p_i^{\text{BR}} - p_j^{\text{BR}}| = \frac{1}{4}|p_j - p_i|$ for any two firms. Since price observations were symmetrically aligned, $|p_j - p_i| < \frac{s}{3}$ holds and therefore $|p_i^{\text{BR}} - p_j^{\text{BR}}| < \frac{s}{3}$ is also satisfied. Thus, adding the best-response prices to the price observations gives a symmetrically aligned set again.³¹

When firms i and j attract both types of consumers while firm k attracts high-type consumers only, then firms learn the following demand parameters: $a_i = a_j = \frac{5}{6}$, $b_{ii} = b_{jj} = \frac{1}{S} + \frac{1}{s}$, $b_{ij} = b_{ji} = \frac{1}{2S} + \frac{1}{s}$, $b_{ik} = b_{jk} = \frac{1}{2S}$, $a_k = \frac{1}{3}$, $b_{kk} = \frac{1}{S}$ and $b_{ki} = b_{kj} = \frac{1}{2S}$ (see Case 2 in the proof of Proposition 2.1 for the corresponding demand functions). Using (3), the best-response prices are given by

$$\begin{aligned} p_i^{\text{BR}} &= \frac{\frac{5}{6} + \left(\frac{1}{2S} + \frac{1}{s}\right)p_j + \frac{1}{2S}p_k}{2\left(\frac{1}{S} + \frac{1}{s}\right)} + \frac{c}{2}, \\ p_j^{\text{BR}} &= \frac{\frac{5}{6} + \left(\frac{1}{2S} + \frac{1}{s}\right)p_i + \frac{1}{2S}p_k}{2\left(\frac{1}{S} + \frac{1}{s}\right)} + \frac{c}{2}, \\ p_k^{\text{BR}} &= \frac{\frac{1}{3} + \frac{1}{2S}(p_i + p_j)}{2\frac{1}{S}} + \frac{c}{2}. \end{aligned}$$

Then $|p_i^{\text{BR}} - p_j^{\text{BR}}| = \frac{\left(\frac{1}{2S} + \frac{1}{s}\right)}{2\left(\frac{1}{S} + \frac{1}{s}\right)}|p_j - p_i| = \frac{2S+s}{4S+4s}|p_j - p_i| < \frac{s}{3}$ since $\frac{2S+s}{4S+4s} < 1$ and $|p_j - p_i| < \frac{s}{3}$ because price observations were asymmetrically aligned. Thus, the first condition in the definition is satisfied.³²

Let us suppose that $p_i \leq p_j$ in the most recent price observation. In that case, $p_j^{\text{BR}} \leq p_i^{\text{BR}}$ and it must hold for having asymmetrically aligned price observations that $p_j^{\text{BR}} + \frac{s}{3} < p_k^{\text{BR}} < p_j^{\text{BR}} + \frac{S}{3}$. Using the formulas above, it can be shown that

$$p_k^{\text{BR}} - p_j^{\text{BR}} = \frac{1}{12(S+s)} [2S^2 - 3Ss - 3Sp_i + (3S+3s)p_j - 3sp_k].$$

We will first show that the condition $p_k^{\text{BR}} - p_j^{\text{BR}} < \frac{S}{3}$ is always satisfied. Using the formula for $p_k^{\text{BR}} - p_j^{\text{BR}}$, the condition simplifies to

$$-\frac{S}{3} \left(2\frac{S}{s} + 7\right) + \left(1 + \frac{S}{s}\right)(p_j - p_i) < p_k - p_i.$$

The left-hand side is smaller than $-\frac{S}{3} \left(2\frac{S}{s} + 7\right) + \left(1 + \frac{S}{s}\right)\frac{s}{3}$ since $p_j - p_i < \frac{s}{3}$. It is easy to see that this expression is always negative. On the other hand, the right-hand side is positive since $p_k - p_i > \frac{s}{3}$. Thus, $p_k^{\text{BR}} - p_j^{\text{BR}} < \frac{S}{3}$ is always satisfied.

³¹Notice the contraction mapping feature of playing the best-response price. This implies that symmetrically aligned prices converge to the same value. Since prices are best response to each other, firms will reach the Nash equilibrium in this case.

³²Note the contraction mapping feature again, which implies that the low-price firms will reach the same price if the information set always remains asymmetrically aligned.

Next let us consider the condition $p_k^{\text{BR}} - p_j^{\text{BR}} > \frac{s}{3}$. Using the formula for $p_k^{\text{BR}} - p_j^{\text{BR}}$, the condition simplifies to

$$\frac{s}{3} \left[2 \left(\frac{S}{s} \right)^2 - 7 \frac{S}{s} - 4 \right] + \left(1 + \frac{S}{s} \right) (p_j - p_i) > p_k - p_i. \quad (\text{A.25})$$

The left-hand side of the inequality is greater than or equal to $\frac{s}{3} \left[2 \left(\frac{S}{s} \right)^2 - 7 \frac{S}{s} - 4 \right]$ as $p_j - p_i \geq 0$. The right-hand side is smaller than $\frac{S}{3}$ since price observations are asymmetrically aligned. Thus, a sufficient condition for (A.25) to hold is that

$$\frac{s}{3} \left[2 \left(\frac{S}{s} \right)^2 - 7 \frac{S}{s} - 4 \right] \geq \frac{S}{3}.$$

This leads to $2 \left(\frac{S}{s} \right)^2 - 8 \frac{S}{s} - 4 \geq 0$, for which $\frac{S}{s} \geq 2 + \sqrt{6}$ must hold. Thus, when the latter condition holds, asymmetrically aligned price observations always remain asymmetrically aligned, irrespective of the exact values in the last price observation. On the other hand, when $\frac{S}{s} < 2 + \sqrt{6}$, condition (A.25) has to hold for the most recent price observation in order to have asymmetrically aligned price observations again.

Since price observations were asymmetrically aligned, $\frac{s}{3} < p_k - p_i$. Thus, the following condition must hold

$$\frac{s}{3} < \frac{s}{3} \left[2 \left(\frac{S}{s} \right)^2 - 7 \frac{S}{s} - 4 \right] + \left(1 + \frac{S}{s} \right) (p_j - p_i).$$

As we have seen, playing the best response works as a contraction mapping for the low-price firms, therefore $p_j - p_i \rightarrow 0$ if the information set always remains asymmetrically aligned. Thus, $\frac{s}{3} \leq \frac{s}{3} \left[2 \left(\frac{S}{s} \right)^2 - 7 \frac{S}{s} - 4 \right]$ must hold. This leads to $1 \leq 2 \left(\frac{S}{s} \right)^2 - 7 \frac{S}{s} - 4$, from which $\frac{S}{s} \geq \frac{7+\sqrt{89}}{4}$. Thus, when the latter condition does not hold, then an asymmetrically aligned information set cannot stay asymmetrically aligned. \square

A.6 The proof of Proposition 3.7

PROOF.

We will now show that both the Nash equilibrium and the ALE are locally stable equilibria. First we will describe the system in the neighborhood of the equilibria and then we show that the eigenvalues of the Jacobian are always less than 1 in absolute value. First we focus on the Nash equilibrium.

Part 1: Stability of the Nash equilibrium

When prices in the information set are symmetrically aligned, then firms learn the correct demand parameters of the linear part on which they operate. Moreover, as we have seen in Lemma 3.6, updating the information set with the best-response prices results in a symmetrically aligned information set again.

Thus, the parameters of the perceived demand functions do not change in this case. Then the perceived demand function of firm i is always given by $D_i^P(p) = \frac{2}{3} - (\frac{1}{S} + \frac{1}{s})p_i + (\frac{1}{2S} + \frac{1}{2s})p_j + (\frac{1}{2S} + \frac{1}{2s})p_k$ (see Case 1 in the proof of Proposition 2.1 for the demand parameters of the relevant linear part). Then the next-period price of firm i is given by

$$p_{i,t+1} = \frac{1}{3} \frac{Ss}{S+s} + \frac{1}{4} (p_{j,t} + p_{k,t}) + \frac{1}{2}c.$$

This holds for every firm i , therefore the Jacobian of the system is given by

$$J = \begin{pmatrix} 0 & \frac{1}{4} & \frac{1}{4} \\ \frac{1}{4} & 0 & \frac{1}{4} \\ \frac{1}{4} & \frac{1}{4} & 0 \end{pmatrix}.$$

The characteristic equation is given by

$$k(\lambda) = -\lambda^3 + \frac{3}{16}\lambda + \frac{1}{32} = 0.$$

It is easy to see that $k(\lambda)$ can also be expressed as $k(\lambda) = -(\lambda - \frac{1}{2})(\lambda + \frac{1}{4})^2$. Thus, the eigenvalues of the Jacobian are $\lambda_1 = \frac{1}{2}$ and $\lambda_2 = -\frac{1}{4}$. Both eigenvalues are smaller than 1 in absolute value, therefore the Nash equilibrium is locally stable.

Part 2: Stability of the ALE

When prices in the information set are asymmetrically aligned, then firms learn the correct demand parameters of the linear part on which they operate. Moreover, as we have seen in Lemma 3.6, updating the information set with the best-response prices results in an asymmetrically aligned information set again when $\frac{S}{s} > \Sigma_2$.³³ Thus, the parameters of the perceived demand functions do not change in this case. Suppose that firms i and j are the low-price firms and firm k is the high-price firm. Then the perceived demand function of firm i is always given by $D_i^P(p) = \frac{5}{6} - (\frac{1}{S} + \frac{1}{s})p_i + (\frac{1}{2S} + \frac{1}{s})p_j + \frac{1}{2S}p_k$ while that of firm k is $D_k^P(p) = \frac{1}{3} - \frac{1}{S}p_k + \frac{1}{2S}p_i + \frac{1}{2S}p_j$ (see Case 2 in the proof of Proposition 2.1 for the demand parameters of the relevant linear part). Then the next-period price of firms i and k are given by

$$\begin{aligned} p_{i,t+1} &= \frac{5}{12} \frac{Ss}{S+s} + \frac{1}{4} \frac{2S+s}{S+s} p_{j,t} + \frac{1}{4} \frac{s}{S+s} p_{k,t} + \frac{1}{2}c, \\ p_{k,t+1} &= \frac{1}{6}S + \frac{1}{4} (p_{i,t} + p_{j,t}) + \frac{1}{2}c. \end{aligned}$$

³³Even if $\frac{S}{s} > \Sigma_2$ does not hold, we can consider a sufficiently small neighborhood of the ALE for which the updated information set is asymmetrically aligned. This can be done as (A.25) holds for (p_L, p_L, p_H) whenever the ALE exists.

The next-period price of firm j is given by a similar formula as for firm i , we just need to switch i and j . Then the Jacobian of the system is given by

$$J = \begin{pmatrix} 0 & A & B \\ A & 0 & B \\ C & C & 0 \end{pmatrix},$$

where $A = \frac{1}{4} \frac{2S+s}{S+s}$, $B = \frac{1}{4} \frac{s}{S+s}$ and $C = \frac{1}{4}$. The characteristic equation is given by

$$k(\lambda) = -\lambda^3 + (A^2 + 2BC)\lambda + 2ABC = 0.$$

It is easy to see that $k(\lambda)$ can also be expressed as $k(\lambda) = -(\lambda + A)(\lambda^2 - A\lambda - 2BC)$. Thus, one eigenvalue is $\lambda_1 = -A$. This eigenvalue is always smaller than 1 in absolute value. $A > 0$ since $S, s > 0$. $A < 1$ if and only if $2S + s < 4(S + s)$, which is always satisfied.

The other two eigenvalues are the solutions of the equation $\lambda^2 - A\lambda - 2BC = 0$. The discriminant is $D = A^2 + 8BC > 0$, so there are two real roots: $\lambda_{2,3} = \frac{A \pm \sqrt{A^2 + 8BC}}{2}$. Root $\lambda_2 = \frac{A + \sqrt{A^2 + 8BC}}{2}$ has the larger absolute value. Its absolute value is smaller than 1 if and only if $\sqrt{A^2 + 8BC} < 2 - A$, from which $A^2 + 8BC < 4 - 4A + A^2$ (using that $A < 1$). This simplifies to the condition $A + 2BC < 1$.

Plugging in the values for A , B and C yields $A + 2BC = \frac{1}{4} \frac{2S+s}{S+s} + 2 \frac{1}{4} \frac{s}{S+s} \frac{1}{4} = \frac{1}{8} \frac{4S+3s}{S+s}$. This is smaller than 1 in absolute value if and only if $4S + 3s < 8(S + s)$, which is satisfied for any $S, s > 0$.

Thus, all three eigenvalues are smaller than 1 in absolute value, implying that the ALE is locally stable.

□

B Robustness check

We report the results of the robustness checks in this section. First we summarize the baseline results in Tables 5 and 6. For each (S, τ) combination we ran 1000 simulations and we investigated how often prices lie in a small neighborhood of the Nash equilibrium or of the ALE. Table 5 shows 4 numbers for each (S, τ) combination. The upper values refer to the Nash equilibrium whereas the lower ones to the ALE. The numbers that are not in brackets correspond to the 0.001-neighborhood of the given equilibrium while the numbers in brackets show the proportion of final price vectors in the 0.0001-neighborhoods.³⁴ As we have discussed before, there is no exact convergence in the model, that is why not all the outcomes lie in the small neighborhood of the equilibria.

³⁴Remember, a vector (x_1, \dots, x_n) lies in the ε -neighborhood of another vector (y_1, \dots, y_n) if their Euclidean distance is smaller than or equal to ε : $\sqrt{\sum_{i=1}^n (x_i - y_i)^2} \leq \varepsilon$.

		τ			
		6	8	10	20
S	2	96.5% (89.7%)	99.8% (97.7%)	100% (99.8%)	100% (100%)
		0.0% (0.0%)	0.0% (0.0%)	0.0% (0.0%)	0.0% (0.0%)
	4	95.0% (87.8%)	99.7% (97.9%)	100% (99.6%)	99.6% (99.6%)
		0.0% (0.0%)	0.0% (0.0%)	0.0% (0.0%)	0.0% (0.0%)
	4.2	59.3% (53.8%)	77.2% (75.6%)	83.6% (83.1%)	88.8% (88.8%)
		37.3% (33.9%)	22.5% (21.9%)	16.4% (15.9%)	11.2% (11.2%)
	4.35	50.8% (46.8%)	73.5% (72.1%)	76.5% (76.2%)	83.3% (83.3%)
		47.4% (43.3%)	26.1% (25.4%)	23.5% (22.9%)	16.5% (16.5%)
	4.5	50.1% (45.5%)	71.4% (70.0%)	72.4% (71.9%)	80.0% (80.0%)
		47.8% (44.0%)	28.3% (26.8%)	27.6% (27.1%)	19.8% (19.8%)
	10	7.3% (7.1%)	12.1% (11.7%)	15.9% (15.9%)	21.0% (21.0%)
		91.7% (79.5%)	87.6% (81.6%)	84.0% (80.9%)	78.2% (78.0%)

Table 5: Proportion of outcomes in the 0.001 and the 0.0001-neighborhoods (in brackets) of the Nash equilibrium (upper numbers) and the asymmetric learning-equilibrium (lower numbers) over 1000 simulations, for different values of S and τ . Other parameters: $s = 1$ and $c = 1$.

	mean	median	mode
random price rule	2 – 13.5	1 – 9	0 – 4
no jump rule	0.35 – 1.21	0 – 1	0
imp. est. rule	0 – $6 \cdot 10^{-3}$	0	0

Table 6: Summary statistics of the usage of augmenting rules.

Table 6 reports information about the usage of the different augmenting rules discussed in Section 3.1. For each (S, τ) combinations we considered, we counted how many times each rule was used during a given initialization and then we calculated the average, the median and the mode of these numbers over the 1000 different initializations. The table reports the range of these statistics over the different (S, τ) combinations. Note that each rule could have been used $996 \cdot 3 = 2988$ times as each firm could use each rule in each period except for the initial ones.

Next we report the results that exclude the effect of the numerical issue that leads to imprecise parameter estimates. Panel *a* of Table 7 summarizes the outcomes of simulations when firms use the true linear

		τ			
		6	8	10	20
S	2	100%	100%	100%	100%
		0%	0%	0%	0%
	4	100%	100%	100%	100%
		0%	0%	0%	0%
	4.2	92.3%	94.4%	94.7%	96.1%
		7.7%	5.6%	5.3%	3.9%
	4.35	90.6%	92.4%	94.8%	94.7%
		9.4%	7.6%	5.2%	5.3%
	4.5	88.0%	90.8%	91.8%	92.6%
		12.0%	9.2%	8.2%	7.4%
	10	41.5%	46.0%	47.0%	49.4%
		58.5%	54.0%	53.0%	49.6%

(a) True coefficients

		τ
		8
S	2	100% (89.1%) 0.0% (0.0%)
	4	100% (81.1%) 0.0% (0.0%)
	4.2	97.5% (77.9%) 2.4% (1.0%)
	4.35	92.7% (74.8%) 7.3% (3.5%)
	4.5	90.4% (72.7%) 9.6% (4.8%)
	10	46.2% (34.4%) 52.2% (13.8%)

(b) Noise

Table 7: Percentage of outcomes in small neighborhoods of the Nash equilibrium (upper numbers) and the asymmetric learning-equilibrium (lower numbers) under the true coefficients (panel a) and under noise (panel b). Number of simulations: 1000, other parameters: $s = 1$ and $c = 1$.

parameter values instead of parameter estimates whenever their information set becomes symmetrically or asymmetrically aligned. Since now there is exact convergence, we show the values that correspond to the 0.0001-neighborhoods only.³⁵ Panel *b* show the results under noise ($N(0, 0.001^2)$). Since noise is added to the model, we consider bigger neighborhoods of the equilibria than in Table 5: The reported numbers correspond to the 0.01 and 0.001-neighborhoods (in brackets) respectively. Notice that the proportions for the 0.01 neighborhoods are quite close to those under the true coefficients.

Finally, Table 8 summarizes the results of simulations with 10000 runs and with 10000 periods for the case $\tau = 8$.

³⁵The values for $(S, \tau) = (10, 20)$ do not add up to 100%. This is caused by the no jump rule: In some simulations one firm got stuck with its last randomly drawn price as the best-response price was too high compared to their most recent price and thus the no jump rule prevented the price change. In these simulations firms reached neither the Nash equilibrium, nor the ALE. However, if we turn off the no jump rule, then such situation does not occur.

		estimated coefficients		true coefficients	
		10000 periods	10000 runs	10000 periods	10000 runs
S	2	99.6% (97.7%)	99.5% (97.9%)	100%	100%
		0.0% (0.0%)	0.0% (0.0%)	0.0%	0.0%
	4	99.5% (97.7%)	99.4% (97.5%)	100%	100%
		0.0% (0.0%)	0.0% (0.0%)	0.0%	0.0%
	4.2	71.5% (70.1%)	78.5% (77.0%)	94.4%	94.8%
		27.8% (26.7%)	21.1% (20.5%)	5.6%	5.2%
	4.35	68.2% (66.6%)	72.1% (70.5%)	92.4%	92.9%
		31.7% (30.9%)	27.6% (26.6%)	7.6%	7.1%
	4.5	62.6% (61.3%)	68.8% (67.5%)	90.8%	91.4%
		37.3% (36.0%)	30.8% (30.0%)	9.2%	8.6%
	10	10.8% (10.6%)	12.4% (12.2%)	46.0%	44.8%
		88.9% (82.9%)	81.4% (81.1%)	54.0%	55.2%

Table 8: Outcome of simulations with 10000 periods and 10000 different initializations. The proportion of outcomes in the 0.001 and the 0.0001-neighborhoods (in brackets) of the Nash equilibrium (upper numbers) and the asymmetric learning-equilibrium (lower numbers) for different values of S and $\tau = 8$. Other parameters: $s = 1$ and $c = 1$.