



The University of Adelaide
School of Economics

Research Paper No. 2013-16
October 2013

Explaining Price Dispersion and Dynamics in Laboratory Bertrand Markets

Ralph-C. Bayer, Hang Wu, and Mickey Chan



Explaining price dispersion and dynamics in laboratory Bertrand markets*

Ralph-C Bayer[†]& Hang Wu[‡]& Mickey Chan[§]

Abstract

This paper develops a quantal-response adaptive learning model which combines sellers' bounded rationality with adaptive belief learning in order to explain price dispersion and dynamics in laboratory Bertrand markets with perfect information. In the model, sellers hold beliefs about their opponents' strategies and play quantal best responses to these beliefs. After each period, sellers update their beliefs based on the information learned from previous play. Maximum likelihood estimation suggests that when sellers have full past price information, the learning model explains price dispersion within periods and the dynamics across periods. The fit is particularly good if one allows for sellers being risk averse. In contrast, Quantal Response Equilibrium does not organize the data well.

JEL Numbers: C73, C91, D83, L13.

Keywords: Price dispersion, Adaptive Learning, Bounded rationality, Quantal Response Equilibrium.

*The authors would like to acknowledge the financial support of the Australian Research Council under DP120101831.

[†]University of Adelaide, School of Economics.

[‡]Corresponding author: University of Adelaide, School of Economics, North Terrace, Adelaide, SA 5005, Australia, email: hang.wu@adelaide.edu.au, Tel.: ++61 (0)8 8313 6024.

[§]University of Adelaide, School of Economics.

1 Introduction

For non-economists it is counter-intuitive that in homogeneous product Bertrand markets, in Bertrand-Nash-equilibrium, all firms always should set prices at the marginal cost level and earn zero profits. The Bertrand-Nash equilibrium prediction is counter-intuitive to non-game theorists in the sense that no firm has strong incentives to stick to the equilibrium because there is no cost of unilaterally deviating from it. Further to that, sellers are able to achieve higher expected profits by deviating together from the Nash equilibrium and coordinating on higher prices. Observations from real world and laboratory markets cast further doubts on the appropriateness of the Bertrand-Nash prediction. In both real and experimental markets, prices are found to be dispersed above the marginal cost. Moreover, rather than staying constant over time, empirical price distributions show significant inter-temporal variation.

Since Stigler (1961), numerous search-theoretical models have been developed to resolve the puzzle. These include Salop & Stiglitz (1977), Reinganum (1979), Varian (1980), and Baye & Morgan (2001), to name only a few. By introducing heterogeneity among consumers or sellers in factors such as search costs, production costs or informational frictions, search-theoretic models provide excellent rationales for price dispersion. Repeated-game effects are also frequently used to explain why prices commonly stay above marginal costs. When sellers engage in repeated interaction, they may have incentives to keep prices at a high level to prevent pricing wars, which might be triggered by a lower price if the competitors adopt a trigger strategy. Alternatively, setting high prices can serve as a signal of friendliness in expectation of reciprocal cooperative behavior by their opponents in later interactions.

Persistent price dispersion, however, is still commonly recorded in markets where all the aforementioned factors play trivial roles. Evidence can be easily found for online-shopping markets, where search and information costs are negligible (Baye & Morgan 2004) . The same is true for laboratory Bertrand markets where all of those factors can be appropriately controlled for (Bayer & Ke 2011; Dufwenberg & Gneezy 2000).

To bridge this gap between theory and empirical reality, some behavioural models have been developed. Rauh (2001) shows that price dispersion can arise when sellers make small but heterogeneous mistakes in beliefs about the market price distribution. Baye & Morgan (2004) show that bounded rationality choice models, namely, Quantal Response Equilibrium (QRE; cf. McKelvey & Palfrey 1995) and ε -equilibrium (Radner 1980) can explain price dispersion in homogeneous-good pricing games. Under the QRE model, firms play quantal best responses in a manner that choices with higher expected profits are played with higher probabilities. In contrast, in an ε -equilibrium sellers are equally likely to choose any price that yields an expected profit within ε of the profits the optimal price would lead to.

Like most of the search-theoretical models of dispersed prices, the bounded rationality equilibrium approaches used by Baye & Morgan (2004) are static and hence fail to capture important dynamic features of market prices. In both laboratory and real world price competition markets, market prices typically exhibit significant intertemporal variation. For this reason, we propose a learning model that combines sellers' boundedly rational pricing behavior with learning, which allows for meaningful dynamics.

Our model is based on laboratory observations from two repeated homogeneous product Bertrand experiments. The only difference between the two experiments is the amount of information revealed to the players after each period. More specifically, in one experiment which we call the *high information* treatment, after each period, firms are shown their private profits and all sellers' prices posted in that period. In contrast, in the *low information* treatment, sellers were shown only their own profits. In both treatments, prices are persistently dispersed over the price set. A comparison of the two treatments shows that the information structure influenced sellers' choices substantially. Prices move downward much faster in the high information treatment than in the low information treatment.

We combine sellers' bounded rationality and learning in an attempt to explain price dispersion and dynamics observed in the high information treatment. In our model, the information about past market prices acts as the main factor of driving the price

adjustments. Following QRE, we assume that the sellers play quantal best responses to their beliefs about the strategies of their opponents. In an extended model, we increase the flexibility of the model by allowing for different risk preferences. We model the dynamics of the game by a belief learning rule. After each period, based on the previous play, sellers update their beliefs about other sellers' strategies and play quantal best responses to the new beliefs. Our model maintains the assumption in Baye & Morgan (2004) that sellers' beliefs take into account the other players' noisy behavior, which is captured by completely mixed strategies. In QRE, all players' beliefs are consistent with the quantal response choices of their opponents. Rather than considering QRE as an instantaneous result of the game, we conjecture that an equilibrium is a steady state of long-run evolution. We assume that the beliefs of a seller (i.e., a probability distribution over the action space of the opponent) is not necessarily correct and changes according to the learning rule. Following Cheung & Friedman (1997), we use an adaptive learning rule with parameter α measuring how past information is discounted. The learning rule includes Cournot learning ($\alpha = 0$) and fictitious learning ($\alpha = 1$) as extreme cases. When $\alpha \in (0, 1)$, all past interactions affect the beliefs; but the more recent periods receive greater weight.

Note that this model lends itself naturally to one of the treatments (i.e. the full information treatment), while it seems highly inappropriate for the other. Clearly, a model built on sellers learning from past prices in the market place only makes sense if the sellers can observe prices. Consequently, our model provides an appropriate explanation for the price dynamics in Bertrand markets with full information only if it fits well in the high information treatment but at the same time not in the low information treatment. Putting the model to the data of the low information treatment serves as a robustness test. If it were to fit well there, then it could not be ruled out that the potentially good fit in the high information treatment is purely mechanical and results from the number of degrees of freedom in the model.

Maximum likelihood estimates show that the quantal-response adaptive learning model nicely captures the price dispersion and dynamic adjustments observed in the

high-information treatment, while the Quantal Response Equilibrium approach of Baye & Morgan (2004) does not. We also find that sellers conditionally on our model being correct exhibit a reasonable degree of risk-aversion. In contrast, our model does not perform well in the the low-information treatment. We conclude that the quantal-response adaptive learning model, where sellers noisily best-respond to their beliefs, is a good explanation for the price dispersion and dynamics in full-information Bertrand markets.

The remainder of this paper is organized as follows. Section 2 introduces two Bertrand price competition experiments and the corresponding data, which we will use as guidance of our modeling. Section 3 lays out the QRE model of Baye & Morgan (2004) and our quantal-response adaptive learning model for the finitely repeated Bertrand market game. Section 4 uses the experimental data to structurally estimate the parameters of the models, and discusses the results. This Section also conveys a comparison of the goodness of fit for the QRE approach and our learning approach. Section 5 concludes with a discussion of the evolutionary properties of the learning model.

2 The experiments

In this section, we present two samples of experimental data which will be taken as guidance for our learning model. In section 4 we will also use these data to evaluate the appropriateness of the model. We use truncated data from two 30-period Bertrand price-competition experiments, one from Bayer & Ke (2011) and the other from a subsequent experiment. Both were conducted at the Adelaide Laboratory for Experimental Economics (Adlab) at the University of Adelaide. In total, 305 participants participated. The participants were mainly students from the University of Adelaide. They studied for a variety of under and postgraduate degrees. The purpose of the experiments was to investigate the effects of exogenous cost shocks on market price. So both experiments have two phases, with 15 periods for each phase. The only difference between the two phases was the sellers' marginal costs. In the first phase the marginal cost was \$E30 for

all sellers and all treatments. After the 15th period, in different treatments the marginal costs either went up to $\$E50$, was kept unchanged at $\$E30$, or declined to $\$E10$.¹ Note that the cost shock was unanticipated such that it should not have any impact on play in the first-phase periods. However, subjects knew that the experiment would run for 30 periods. Since the effects of exogenous cost shocks are not of interests for this study, we focus on the first 15 periods of play from these experiments for which the production cost kept constant. Our results can therefore be interpreted as valid for initial learning in longer repeated Bertrand games.

At the beginning of the experiments the participants were randomly assigned roles as sellers or buyers at a fixed ratio of two to one. The roles were kept fixed throughout the experiment. In each period, markets were formed using random matching. Each market consisted of two sellers and a buyer and all subjects were assigned to participate in a market. Random re-matching was adopted to minimize repeated game effects. In each market, two sellers simultaneously and independently set integer prices that could range from 30 (marginal cost) to 100 (reservation value for the buyer). Afterwards, the buyer observes both prices costlessly and then chooses either to buy from one of the sellers or to leave without buying.² In each stage the payoff for a seller who managed to sell was her price less the cost. An unsuccessful seller earned a profit of zero. The buyers' payoffs were defined as their reservation value minus the price they paid if they bought and zero if they did not buy. The only difference between the two experiments lies in the information that was revealed to the subjects at the end of each period. In the high information treatment (120 participants), all players were shown their profits and the prices set by *all* sellers in the session³. In the low information treatment (185 participants), the participants learned only their own payoffs. No price information was given. Before the experiments, participants were provided with written instructions containing the market rules and the payoff functions. At the end of the experiments, the

¹The currency was Experimental Dollars. In what follows we drop the currency symbols.

²In more than 99% cases the buyers bought from the seller with the lower price.

³There were between 12 and 18 sellers in a session.

participants were paid according to their aggregate payoffs in the experimental session. On average they earned around 20 Australian Dollars for about one hour of their time.

Insert Figure 1 about here.

Figure 1 shows the time series for the interquartile ranges of the prices (boxes) as well as the average prices (black lines). Red bars in the boxes represent the median price levels. For both high information and low information treatments, prices were dispersed persistently over the price set. As can be seen from Figure 1, in both panels, the central 50 percent of prices exhibit substantial spreads for all periods.

In terms of price dynamics, however, starting at virtually identical distributions the prices developed quite differently between the two treatments. For the high information treatment, the average price started off at 59.8, with an interquartile range of 50 to 70 and a median price at 60. Then the prices declined quickly as the experiment proceeded. In period 15, the average price was 40.5, and median price was 38, with an interquartile range of 34 to 42. For the low information treatment, the prices started off at similar levels as in the high information treatment. In period 1, the mean and median prices were 60.5 and 60, respectively, and the corresponding interquartile range was 53.5 to 65. While the prices kept dropping quickly towards Nash equilibrium in the high information treatment, the prices declined at a much slower speed and stabilized in the low information treatment. In period 15, the average price was 48.8, the median price was 49, and the interquartile range 45 to 51. All of these characteristic values are about 10 units above their counterparts in the high information treatment. The fact that the prices evolve significantly differently in the two treatments suggests that feedback on the past strategies plays an important role in price dynamics. In our learning model, the feedback effect will be considered as the main driving force of the price adjustments.

3 Theory

3.1 Preliminaries

Consider an environment where a set $I = \{1, 2, \dots, N\}$ of sellers engage repeatedly in the standard Bertrand duopoly game along the time horizon $T \equiv \{1, 2, \dots, 15\}$. At the beginning of each period, each seller $i \in I$ is randomly matched with a competing seller $j \in I$. Afterwards, seller i and j compete in prices to sell a homogeneous good produced at cost c per unit. The market has unit demand for the good up to a reservation price v . Without loss of generality, we define the price set as $P \equiv [c, v]$. Let (p_i, p_j) be the prices set by the two competing sellers, the payoff to seller i is

$$\pi_i(p_i, p_j) = \begin{cases} p_i - c & \text{if } p_i < p_j \\ \frac{1}{2}(p_i - c) & \text{if } p_i = p_j ; \forall p_i, p_j \in P. \\ 0 & \text{otherwise} \end{cases} \quad (3.1)$$

Let seller i 's strategy be a cumulative probability measure over the price set, denoted as $F_i : P \rightarrow [0, 1]$. Further, let $B_i(F_j)$ be seller i 's belief about her rival j 's strategy. Thus, the expected monetary payoff for seller i posting price p , given $B_i(F_j)$, is

$$E\pi_i(p) = (p - c) [1 - B_i(F_j(p))], \forall i, j \in I, i \neq j. \quad (3.2)$$

We now state, without proof, the well known Bertrand-Nash equilibrium, where all probability mass in a mixed strategy is put on the price that equals marginal cost.

Proposition 1. (*Bertrand-Nash equilibrium*) *For all periods $t \in T$, the following comprises a symmetric Bertrand-Nash equilibrium: For all $i, j \in I$ and for all $p \in P$, $F_i^{NE}(p) = 1$.*

3.2 Quantal Response Choices

Our first extension to the Nash equilibrium is to introduce bounded rationality to sellers' pricing decisions. We assume that sellers are prone to choice errors and post sub-optimal prices with positive probabilities. The errors could be caused by inexperience, computational limits, or instantaneous mood shocks (see Chen et al. 1997). Following Baye & Morgan (2004) and López-Acevedo (1997), we incorporate the choice errors using a power-form quantal-response function. Formally, the strategy of seller $i \in I$ in terms of cumulative probability distribution is

$$F_i(p, B_i) = \frac{\int_c^p [E\pi_i(q, B_i(F_j))]^\lambda dq}{\int_c^v [E\pi_i(k, B_i(F_j))]^\lambda dk}, \forall p \in P \quad (3.3)$$

where

$$E\pi_i(p, B_i(F_j)) = (p - c) [1 - B_i(F_j(p))]. \quad (3.4)$$

The probability that seller i set her price at p is positively related to the expected monetary payoff resulting from p . The “error parameter” or “bounded-rationality parameter,” $\lambda \in [0, \infty)$, measures the degree of sensitivity of the firms to the expected payoffs. As $\lambda \rightarrow \infty$, the firm tends to choose the payoff maximizing price with certainty and becomes fully rational. On the other hand, as $\lambda \rightarrow 0$, the firm becomes fully ignorant or confused and randomizes over all prices with equal probabilities. A property of the choice-probability functions that is important for empirical applications and also for economic intuition is that strategies that yield greater expected payoffs are chosen with higher probabilities when $\lambda > 0$. A noteworthy special case is $\lambda = 1$, under which choice probabilities of firm i are proportional to the expected payoffs. This is the classic Luce (1959) probabilistic choice model and was first applied in non-cooperative games by Rosenthal (1989).

3.3 Quantal Response Equilibrium (QRE)

In a Quantal Response Equilibrium (QRE), all sellers' strategies are quantal responses to their beliefs about the probability distributions of their opponents' prices. That is, for each $i \in I$, F_i follows Equation (3.3). Moreover, the beliefs of all sellers are consistent with the probability distributions of their opponents' prices. We have $B_i(F_j) = F_j, \forall i, j \in I$. Baye & Morgan (2004) obtained a simple, closed-form representation of QRE pricing strategies for the homogeneous product Bertrand duopoly game.

Proposition 2. (*Baye & Morgan 2004*) *For any $\lambda \in [0, 1)$, the following comprises a symmetric QRE:*

$$F_i^Q(p) = 1 - \left[\frac{\pi(v)^{1+\lambda} - \pi(p)^{1+\lambda}}{\pi(v)^{1+\lambda} - \pi(c)^{1+\lambda}} \right]^{\frac{1}{1-\lambda}} \quad \forall p \in [c, v], \forall i \in I \quad (3.5)$$

where $\pi(p)$ is the payoff to a monopolist charging price p .

Proof. See Baye & Morgan (2004). □

Note that for the QRE to exist we must have $\lambda \in [0, 1)$. When λ is greater than one, according to Equation (3.5), we have $F_i^Q(p) > 1$ for all prices that are between c and v , which is impossible. With $\lambda = 0$, the sellers behave randomly in choosing prices so that the prices are distributed uniformly over the price set P . At the other extreme, with λ tends to 1, more and more probability mass is allocated to low prices and the QRE converges to the Bertrand-Nash equilibrium. Therefore, in this game, to attain the Bertrand-Nash equilibrium result, perfect rationality ($\lambda \rightarrow \infty$) is not required.

An appealing feature of the quantal-response choice rule is its flexibility that allows for incorporating and parametrization of factors that may influence players' behavior

other than bounded rationality. By using the Arrow-Pratt risk measure, we can extend the QRE model to allow for heterogeneous attitudes toward risk or uncertainty in different circumstances of the game.⁴ Formally, instead of maximizing the expected monetary payoffs, we assume that the sellers aim to maximize expected utilities which we define as

$$EU_i(p) = \frac{(p-c)^{1-r}}{1-r} [1 - B_i(F_j(p))]; \forall t \in T, \forall i, j \in I. \quad (3.6)$$

The parameter r measures a seller's risk attitudes, with $r = 0$ corresponding to risk neutrality, $r > 0$ to risk aversion, and $r < 0$ to risk seeking.⁵

Proposition 3. (*QRE with Arrow-Pratt Risk Attitudes*) For any $\lambda \in [0, 1)$ and $r < 1$, the following comprises a symmetric QRE with risk attitudes:

$$F_i^{AP}(p) = 1 - \left[1 - \left(\frac{p-c}{v-c} \right)^{(1-r)(1+\lambda)} \right]^{\frac{1}{1-\lambda}}. \quad (3.7)$$

Proof. Setting $\pi(p) \equiv \frac{(p-c)^{1-r}}{1-r}$ in Equation (3.5) yields the result. □

⁴See Goeree et al. (2002) for an example that incorporates QRE with risk aversion in explaining overbidding in private value auctions.

⁵ The utility function we use exhibits constant relative risk aversion and is used frequently in experimental research (e.g., Holt & Laury 2002). For $r = 1$, where the expected utility function is undefined, we use $\ln(p-c)$ instead of $\frac{(p-c)^{1-r}}{1-r}$. This is because for $r \rightarrow 1$ we have $\frac{d \ln(x)}{dx} = \frac{d(\frac{x^{1-r}}{1-r})}{dx}$.

3.4 A Quantal Response Adaptive Learning (QRAL) model

In this subsection we propose a simple quantal response learning model to explain sellers' intertemporal price adjustments observed in the high information treatment. A Quantal Response Equilibrium has the property that all players play noisy best responses to each others noisy best responses. This can be interpreted as the limiting case of subjects playing noisy best responses to their beliefs, which they update between rounds until beliefs and actual play coincide. In order to model this learning process we assume that firms formulate beliefs of their competitors' future strategies based on the price information of the past periods and play quantal responses to their beliefs. We use an approach similar to the empirical learning rule of Cheung & Friedman (1997) where a player's belief is the weighted average of the strategies that she encountered in the past periods. Cournot learning and fictitious learning are special cases of the model. While Cheung & Friedman (1997) assume that players' beliefs are formulated using the past strategies of their actual rivals, we assume that the sellers' current beliefs are the weighted average of all her potential opponents' past strategies. This is a reasonable assumption because in our context the sellers are randomly rematched in each period and are shown the prices of all sellers in the same session. Let (F_1, \dots, F_t) denote the vector of market price distributions observed from period 1 to period t , the belief firm i holds before period $t + 1$ is:

$$B_{i,t+1}(F_{j,t+1}) = \frac{F_t + \sum_{\tau=1}^{t-1} \alpha^\tau F_{t-\tau}}{1 + \sum_{\tau=1}^{t-1} \alpha^\tau}. \quad (3.8)$$

Parameter α captures the idea that different past histories enter with different weights into the beliefs. When $0 < \alpha < 1$ we have the typical case that recent histories carry more weight than older histories. Setting $\alpha = 0$ yields the Cournot adjustment rule, where only the most recent period is relevant for the beliefs. Setting $\alpha = 1$ yields standard fictitious play, where all past experiences are weighed evenly. Consequently, in period $t + 1$ seller i 's probability choice function, given her belief $B_{i,t+1}(F_{j,t+1})$, can be written as

$$F_{i,t+1}(p) = \frac{\int_c^p [EU_i(q, B_{i,t+1}(F_{j,t+1}))]^\lambda dq}{\int_c^v [EU_i(k, B_{i,t+1}(F_{j,t+1}))]^\lambda dk}. \quad (3.9)$$

4 Estimation

In this section we use our experimental data to estimate the parameters of the QRE and QRAL models and evaluate the relative fit of these models. We use a discretized version of the quantal response function and adopt the following *interiority* condition (cf. Goeree et al. 2005):

$$f_{i,t}(p, B_{i,t}) = \frac{[EU_i(p, B_{i,t}(F_{j,t}))]^\lambda}{\sum_{k=c}^v [EU_i(k, B_{i,t}(F_{j,t}))]^\lambda} > 0, \forall i, j \in I, \forall t \in T, \forall p \in P.$$

That is, the mixed strategies defined by the quantal response functions are complete so that all prices in P are played with positive probability. An example in which the *interiority* condition is violated is $B_i(F_j(p)) = 1$ for all p in P . In this case, player i believes that c is played with certainty by firm j , so the expected payoff for any price is zero. Hence, both the numerator and denominator of the quantal response function are equal to zero, which causes an indeterminacy problem. To avoid such indeterminacy and to ensure that the *interiority* condition is satisfied, in our estimations, we adjust the expected payoffs by adding a small positive technical parameter ε :

$$EU_{i,t}(p, B_{i,t}) = \varepsilon + \frac{(p - c)^{1-r}}{1 - r} [1 - B_{i,t}(F_{j,t}(p))]; \forall t \in T, \forall i, j \in I \quad (4.1)$$

One justification for ε is that when people take part in economic activities, they receive some level of satisfaction from participating, which is independent of the monetary outcomes they get from the activities. For example, in most economics experiments, subjects are rewarded with a show-up fee for participation in addition to earnings that are proportional to their performance. The introduction of ε considerably facilitates the empirical application of the model. Also, when ε is sufficiently small, compared to the general expected payoffs, it will not change any of the main implications of the model. After the transformation, when facing $B_{i,t}(F_{j,t}(p)) = 1, \forall p \in P$, the quantal response choice function will assign uniform probabilities to all prices in P , which is intuitively and economically more convincing because we naturally expect prices with identical payoffs to carry the same weight in the sellers strategies.⁶

We use maximum likelihood estimation (MLE) to derive the estimates for the parameters of interest. To do this, we search numerically for the parameters that maximize the likelihood of occurrence of the set of prices observed in the experiments. We set $\varepsilon = 10^{-10}$ and take the first period's price distribution as the sellers' initial beliefs. From period 2 to 15, for each treatment we calculate the probabilities associated with the prices chosen by sellers using the quantal response function. The log-likelihood function is

$$\log(L) = \log \left[\prod_{t=2}^{15} \prod_{i=1}^N f_{i,t}(p_{i,t}, B_{i,t}) \right] = \sum_{t=2}^{15} \sum_{i=1}^N \log [f_{i,t}(p_{i,t}, B_{i,t})]; \quad (4.2)$$

where N is the number of sellers participating in a treatment. For the purpose of comparison, we conduct the ML estimations for both the QRE and QRAL models in both high and low information treatments. Recall that the low information treatment is

⁶ The parameter ε is required in the power-function specification to prevent indeterminacy. We use the power specification to keep our model in line with Baye & Morgan (2004). Alternatively, we could have used the logistic specification, which allows for zero and negative payoffs.

inconsistent with our learning model as the seller does not have the information required. The fit in this situation, where the model is misspecified by design, will be used as a robustness test.

Insert Table 1 about here.

Table 1 reports the maximum likelihood estimates for the high information treatment. In parentheses are standard errors obtained using numerical differentiation.⁷ The table also includes the log-likelihood value $\log(L)$ and the Bayesian Information Criterion (BIC), which is used to compare the relative goodness of fit for the different models.⁸ According to BIC, QRAL outperforms QRE in a substantial way. Adding a risk-preference parameter improves the fit even when the additional degree of freedom is taken into account. In the learning model with risk preference parameter, the decay parameter $\hat{\alpha}$ is estimated as 0.505, which is significantly different from both zero and one. This is reasonable because we would naturally expect the influence of past history to decay as the experiment proceeds. With $\hat{\alpha} = 0.5$, the price information from more than four periods ago has lost almost all of its influence on today's beliefs (as $0.5^5 \approx 0.03$).

For both QRE and QRAL models, the parameter \hat{r} is positive, which indicates that sellers' choices were guided by risk aversion.⁹ For the QRE models, the estimates for the bounded rationality parameter $\hat{\lambda}$ are equal to 0.914 and 0.842, respectively for the

⁷We have also estimated two alternative models as robustness checks. We first ran the estimation using a truncated data set (from period 2 to 12) and got $\hat{r} = 0.242$, $\hat{\lambda} = 1.567$, $\hat{\alpha} = 0.480$. Secondly, we estimated the parameters for 15 periods but used a uniform distribution as the initial beliefs and obtained $\hat{r} = 0.417$, $\hat{\lambda} = 1.273$, $\hat{\alpha} = 0.412$.

⁸In our analysis, BIC is defined as $BIC = \frac{k}{2} \ln(N * T) - \ln(L)$. Here k is the number of parameters, N is the number of sellers, T is the number of periods considered, and $\ln(L)$ the value of the log-likelihood function. BIC penalizes models with additional parameters. According to this criterion, a model with a lower BIC value is preferred.

⁹When $\hat{r} = 0.365$, a subject is willing to pay about 35 dollars to take a gamble that yields zero and 100 with the same probability 0.5.

models with and without risk parameter. In contrast, if we allow for learning, both $\hat{\lambda}$'s are about 1.5. This suggests that with QRAL the price dynamics may fail to converge to a QRE because for QRE to exist, as indicated in subsection (3.3), λ needs to be less than one. So far we can conclude that a model that combines a belief learning with quantal best response behavior dominates the static QRE model typically used to explain pricing behavior in Bertrand duopolies. Adding a risk parameter further improves the explanatory power of the model.

Insert Figure 2 about here.

Figure 2 plots the estimated mean prices for the adaptive learning models, along with the empirical mean prices for the high-information treatment. It also shows the mean prices of simulations using the estimated parameters. In the simulations, we adopt the market-price distribution of period one as the initial belief, $B_{i,2}(F_{j,2}) = F_1$. Then given the values of \hat{r} , $\hat{\alpha}$ and $\hat{\lambda}$, we can obtain the predicted mixed strategies for period two. Then instead of using the actual observed price distributions, the simulations use the predicted strategies to form the new beliefs and proceed by iterating forward on the system to obtain the simulated strategies for all periods. Therefore, the difference between the simulations and estimations is that in the simulations we use the strategies predicted by Equation (3.9) to formulate the beliefs, while in the estimations we use actually documented strategies. As can be seen from Figure 2, the average prices predicted by the QRAL model are fairly close to the empirically observed dynamics. In particular, in our preferred model (i.e. QRAL with risk preferences) all three time series – empirical, estimated and simulated prices – are very close together.

Insert Figure 3 about here.

Figure 3 shows the estimated distributions of prices predicted by the QRE and QRAL models separately for the 15 periods, both with risk-preference parameters, along with

the empirical price distributions observed in the experiments. The plots indicate that the QRAL model predicts the price dispersion and its evolution quite well, while the QRE model, due to its static nature, works only well for some middle periods and fails to capture the price distribution dynamics.

Insert Table 2 about here.

Insert Figure 4 about here.

When we apply the same methodology to the low information treatment data, we find that the estimated parameter values are not plausible and the model fit is poor. Table 2 shows our estimation results for the different models. In order to achieve a reasonable fit we require an unreasonably low risk-preference parameter r of less than -0.6 . In the otherwise best-fitting QRAL model, this implies an unreasonably high level of risk-love.¹⁰ Moreover, it becomes clear that, when we plot estimated and simulated price time series against the observed prices (Figure 4), that the model does a poor job at explaining the pricing behavior in the low-information treatment. The poor performance of the model in the low information treatment indicates that the model's good fit in the high information treatment is not merely an artifact of its degrees of freedom. Consequently, we conclude that the estimated QRAL model with mildly risk-averse subjects is a robust explanation for observed behavior in Bertrand duopolies with perfect information on past prices.

5 Concluding remarks

In this paper we developed a quantal-response adaptive learning model in order to

¹⁰A person with such risk preferences prefers a gamble of \$2 with probability 1/3 and nothing with probability 2/3 to receiving \$1 for certain.

explain price dispersion and dynamic adjustments observed in repeated experimental Bertrand markets, where prices are observable. In our model, rather than being fully rational and choosing only crisp best responses, the sellers are assumed to be boundedly rational in the sense that they play suboptimal strategies with positive probabilities. The probability they play a specific strategy with is a monotonic function of that strategy's expected payoff. We show that price dispersion can be effectively explained by such quantal-response choice rules. However, the static equilibrium approach based on the quantal-response choice rule, QRE, fails to explain the evolution of prices over time. We use an adaptive belief-learning rule to model learning and to explain price dynamics. The beliefs of the sellers are assumed to be the weighted average of her rivals' past strategies. We show that for experiments, where the sellers have perfect information on their opponents' past choices, the quantal-response adaptive learning model can explain the dynamic evolution of prices remarkably well. In contrast, the model fails to provide reasonable estimates for experiments where no price information of preceding play was revealed. This result has two important implications. Firstly, the good fit of the adaptive learning model is not an artifact just stemming from the model's degrees of freedom. Secondly, an alternative model with limited past information is necessary for the low-information treatment.

We want to conclude with an out-of sample investigation. It is interesting to study how the price dynamics of our best model will evolve if we extend our investigation to a time horizon that is longer than what we had in the experiments. We simulate the QRAL model using the estimated parameters for 200 periods. The result shows that the average price will evolve cyclically without stabilizing at an equilibrium. Figure 5 shows the evolution of the average price. The intuition behind the cyclicity is the following: whenever the mass of the price distribution gets pushed towards marginal cost, then the profitability of charging a price close to marginal cost is very low. thus it becomes profitable to charge a higher price and hope for the rare occurrence of a competitor who charges a high price due to bounded rationality. Given the adaptive nature of the learning process, many sellers will follow this strategy at the same time. Thereafter

the downwards dynamics sets in again until a jump becomes profitable again. In our experiments, we do not have enough time periods in order to see if the cyclical pattern emerges. Bruttel (2009) found some cyclical movement in her series of experiments, which provides some evidence.

Insert Figure 5 about here.

Theoretical explanations of price cycles have focused on Edgeworth cycles. Maskin & Tirole (1988) show that in dynamic Bertrand duopoly games, if sellers follow alternating-move dynamics and adopt Markov perfect equilibrium, then cyclical prices will be a natural result. In their model, sellers engage in price undercutting until they arrive at a bottom price, at which the equilibrium strategy for the firm who gets to move is to raise prices with positive probability. When the firm raises its price, a new price cycle is triggered. Our model provides an alternative explanation for the cyclical price phenomenon in Bertrand markets. As opposed to the alternating-move assumption that only one firm gets to move in each period, we allow both firms to adjust their prices in all periods. Chen et al. (1997) show analytically that for any finite game where the payoffs are positive for all players, if the choices of players are noisy enough, or put equivalently, if the bounded rationality parameter λ is small enough, then fictitious play converges to a unique Quantal Response Equilibrium. For our adaptive learning model, conditions in which the prices converge and in which the dynamics fail to converge still need to be investigated formally. Moreover, a more in-depth experimental investigation is required in order to test if price-cycles occurring in laboratory studies are consistent with our theory.

References

- Baye, M. R. & Morgan, J. (2001), ‘Information gatekeepers on the internet and the competitiveness of homogeneous product markets’, *American Economic Review* **91**, 454–474.
- Baye, M. R. & Morgan, J. (2004), ‘Price dispersion in the lab and on the internet: Theory and evidence’, *The RAND Journal of Economics* **35**(3), 449–466.
- Bayer, R. C. & Ke, C. (2011), ‘Are rockets and feathers caused by search or informational frictions?’, *Working Papers* .
- Bruttel, L. (2009), ‘Group dynamics in experimental studies; the bertrand paradox revisited’, *Journal of Economic Behavior & Organization* **69**(1), 51–63.
- Chen, H. C., Friedman, J. W. & Thisse, J. F. (1997), ‘Boundedly rational nash equilibrium: a probabilistic choice approach’, *Games and Economic Behavior* **18**(1), 32–54.
- Cheung, Y. & Friedman, D. (1997), ‘Individual learning in normal form games: Some laboratory results’, *Games and Economic Behavior* **19**(1), 46–76.
- Dufwenberg, M. & Gneezy, U. (2000), ‘Price competition and market concentration: an experimental study’, *International Journal of Industrial Organization* **18**(1), 7–22.
- Goeree, J. K., Holt, C. A. & Palfrey, T. R. (2002), ‘Quantal response equilibrium and overbidding in private-value auctions’, *Journal of Economic Theory* **104**(1), 247–272.
- Goeree, J. K., Holt, C. A. & Palfrey, T. R. (2005), ‘Regular quantal response equilibrium’, *Experimental Economics* **8**(4), 347–367.
- Holt, C. & Laury, S. (2002), ‘Risk aversion and incentive effects’, *The American economic review* **92**(5), 1644–1655.

- López-Acevedo, G. (1997), ‘Quantal response equilibria for posted offer-markets’, *Estudios Económicos* pp. 95–131.
- Luce, R. (1959), *Individual choice behavior.*, John Wiley.
- Maskin, E. & Tirole, J. (1988), ‘A theory of dynamic oligopoly, ii: Price competition, kinked demand curves, and edgeworth cycles’, *Econometrica: Journal of the Econometric Society* pp. 571–599.
- McKelvey, R. D. & Palfrey, T. R. (1995), ‘Quantal response equilibria for normal form games’, *Games and Economic Behavior* **10**(1), 6–38.
- Radner, R. (1980), ‘Collusive behavior in noncooperative epsilon-equilibria of oligopolies with long but finite lives’, *Journal of Economic Theory* **22**(2), 136–154.
- Rauh, M. (2001), ‘Heterogeneous beliefs, price dispersion, and welfare-improving price controls’, *Economic Theory* **18**(3), 577–603.
- Reinganum, J. F. (1979), ‘A simple model of equilibrium price dispersion’, *The Journal of Political Economy* pp. 851–858.
- Rosenthal, R. W. (1989), ‘A bounded-rationality approach to the study of noncooperative games’, *International Journal of Game Theory* **18**(3), 273–292.
- Salop, S. & Stiglitz, J. (1977), ‘Bargains and ripoffs: A model of monopolistically competitive price dispersion’, *The Review of Economic Studies* **44**(3), 493–510.
- Stigler, G. J. (1961), ‘The economics of information’, *The Journal of Political Economy* **69**(3), 213–225.
- Varian, H. R. (1980), ‘A model of sales’, *The American Economic Review* **70**(4), 651–659.

	risk neutral($r = 0$)				risk seeking/averse ($r \neq 0$)				
	λ	α	$\log(L)$	BIC	λ	α	r	$\log(L)$	BIC
QRAL ($\alpha \in [0, 1]$)	1.542 (0.021)	0.321 (0.010)	-3883	3890	1.522 (0.014)	0.505 (0.029)	0.365 (0.017)	-3826	3837
QRE	0.914 (0.022)	–	-4122	4125	0.842 (0.011)	–	0.255 (0.003)	-4027	4034

Table 1: Maximum Likelihood Estimates for the High Information Treatment.

	risk neutral($r = 0$)				risk seeking/averse ($r \neq 0$)				
	λ	α	$\log(L)$	BIC	λ	α	r	$\log(L)$	BIC
QRAL ($\alpha \in [0, 1]$)	1.578 (0.021)	0.865 (0.016)	-6176	6183	1.436 (0.016)	0.476 (0.009)	-0.602 (0.010)	-6101	6112
QRE	0.836 (0.011)	–	-6411	6415	0.894 (0.021)	–	-0.235 (0.011)	-6347	6354

Table 2: Maximum Likelihood Estimates for the Low Information Treatment.

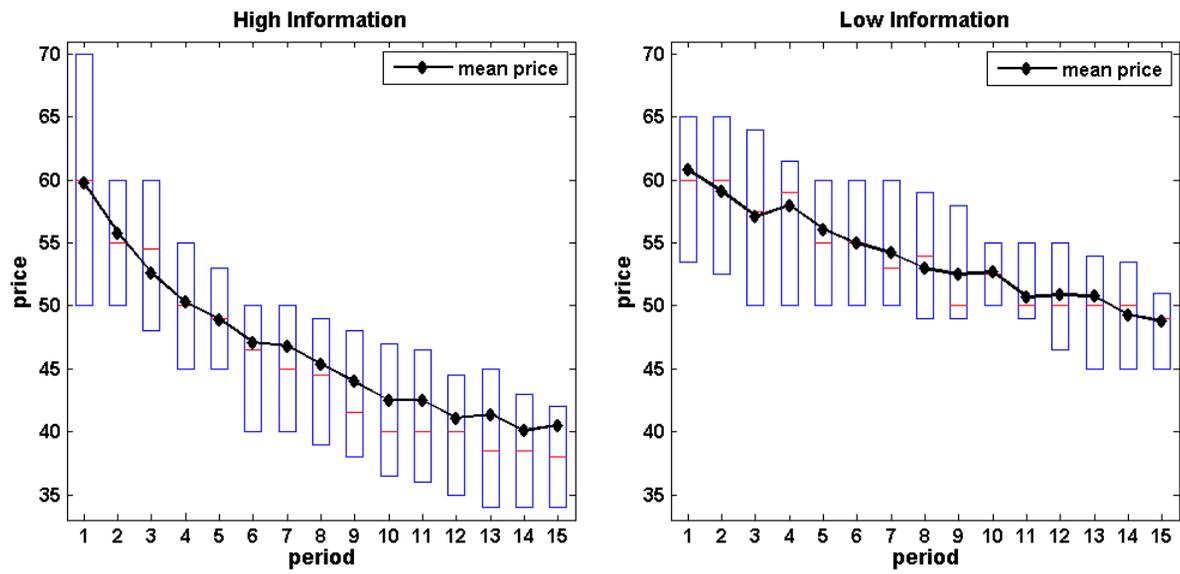


Figure 1: Time Series for Interquartile Ranges of Prices and Mean Prices

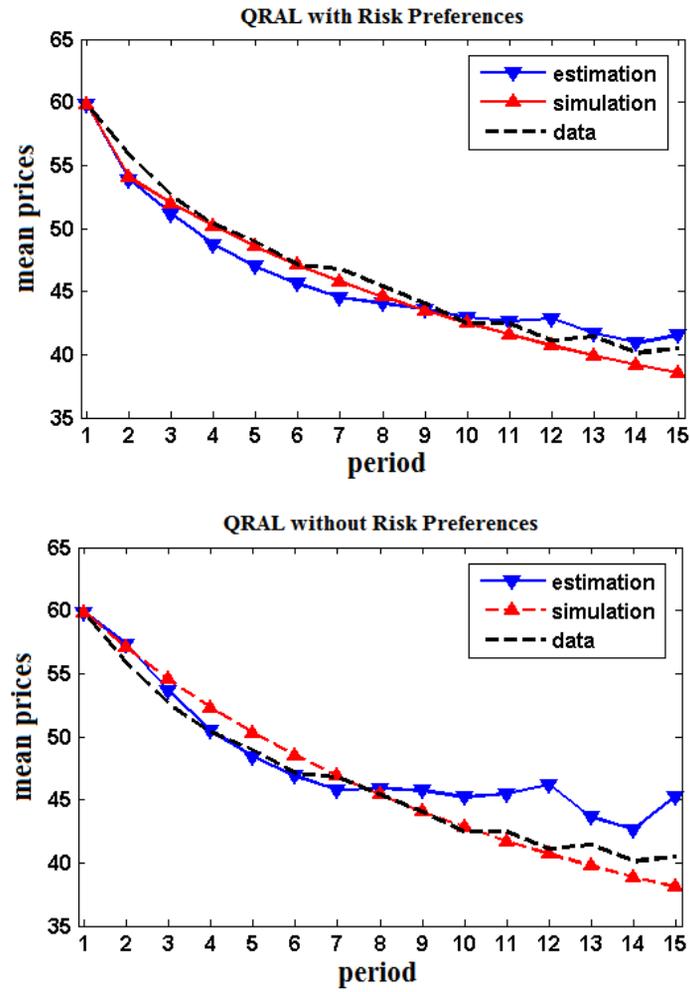
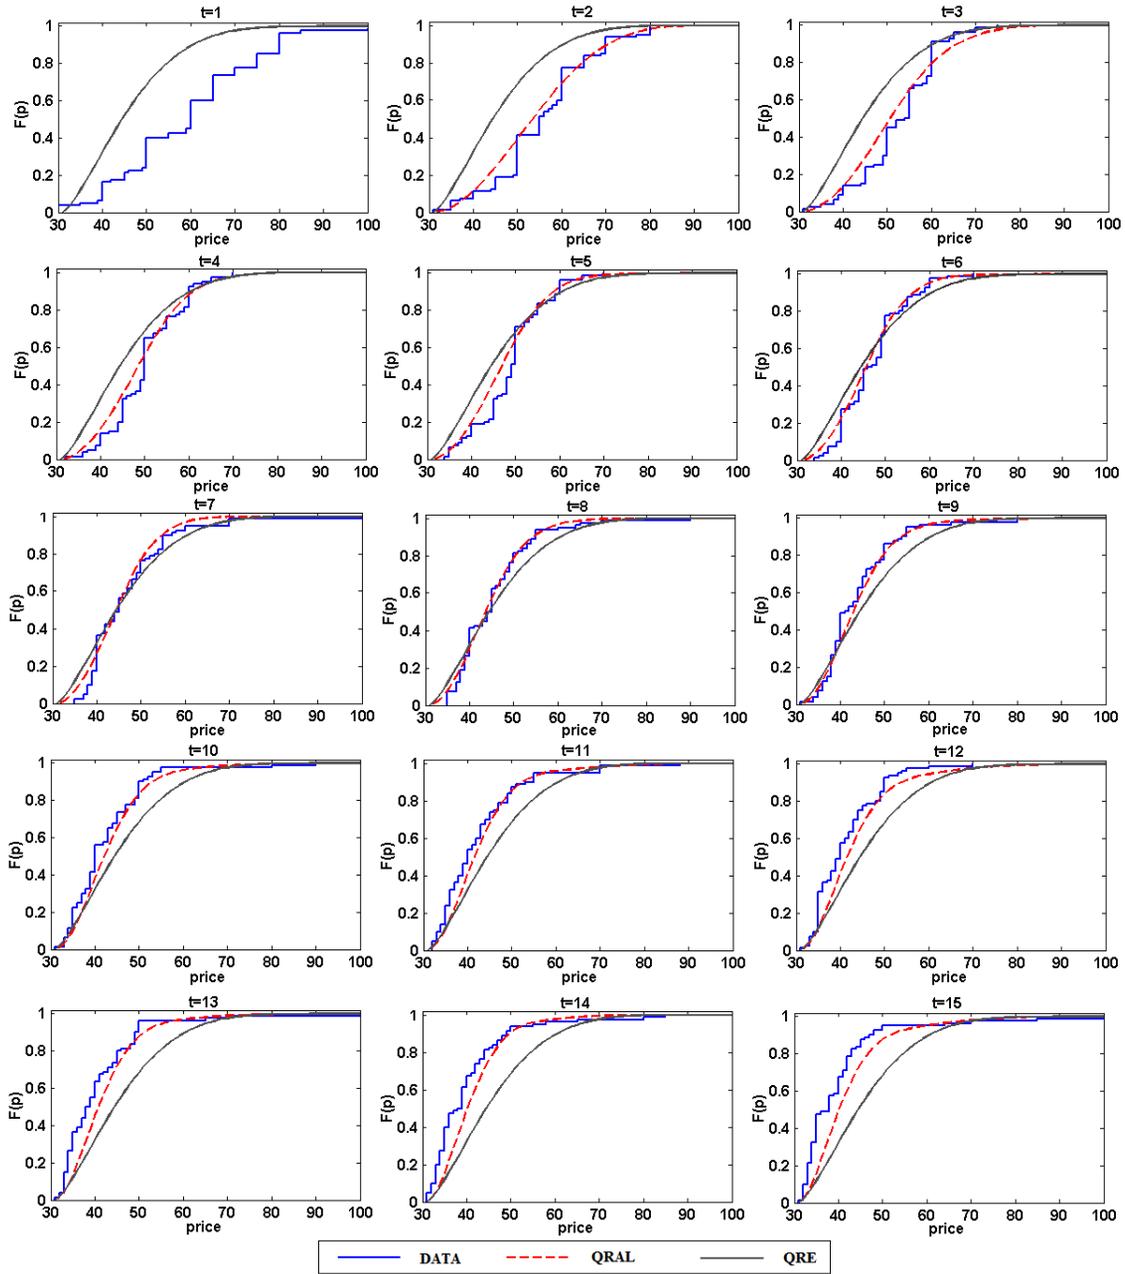


Figure 2: Time Series of Mean Prices: Data and QRAL Predictions-High Information Treatment



yie

Figure 3: Comparisons of Cumulative Distributions of Prices: Data, QRAL with Risk Preferences, and QRE with Risk Preferences

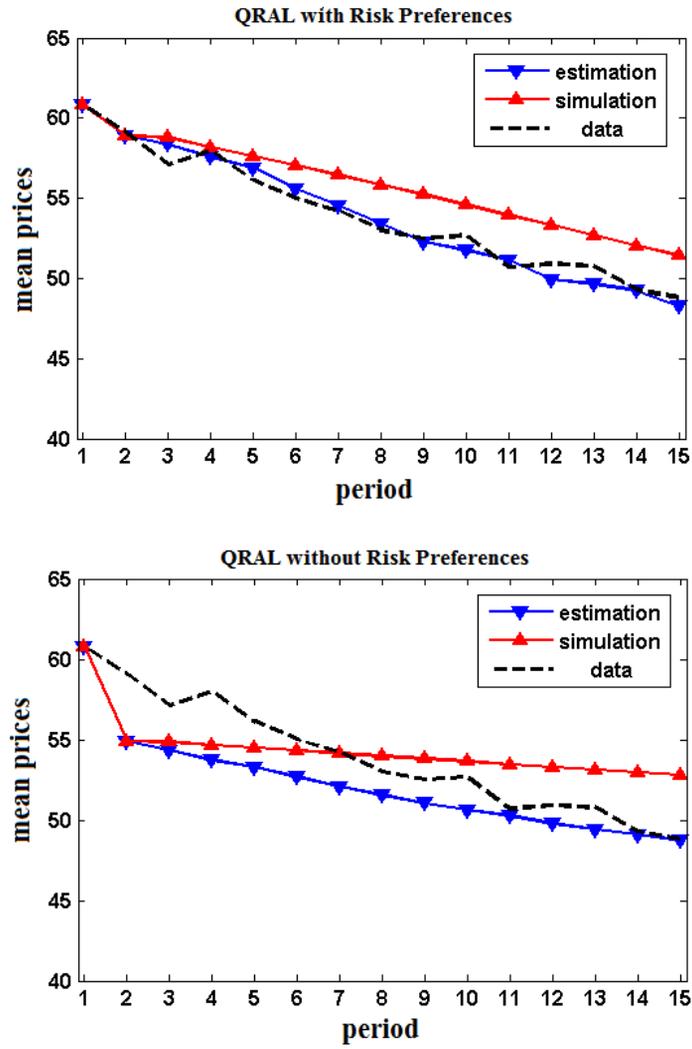


Figure 4: Time Series of Mean Prices:Data and QRAL Predictions-Low Information Treatment

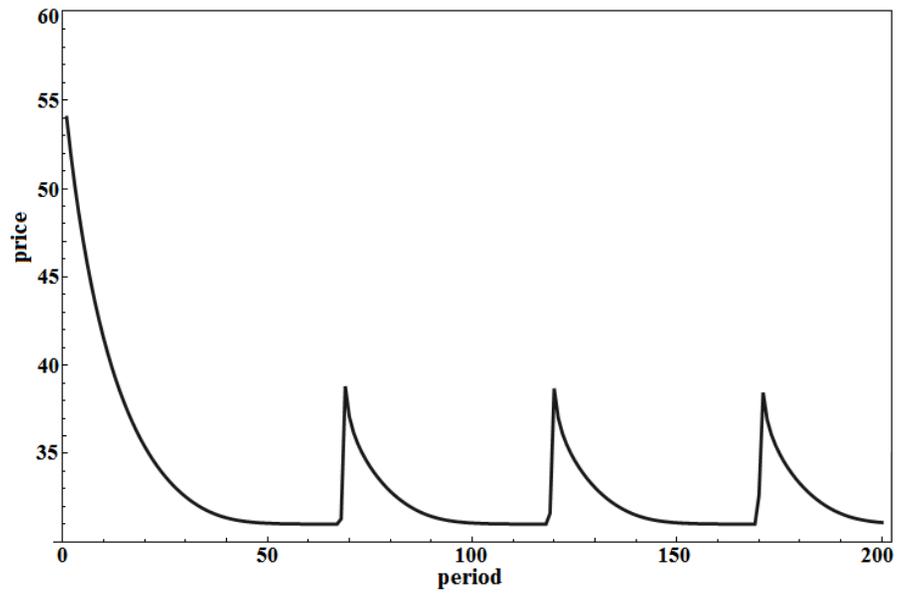


Figure 5: Simulation of the QRAL model with estimated parameters for 200 periods