



The University of Adelaide
School of Economics

Research Paper No. 2010-20
October 2010

Rockets and Feathers in the Laboratory

Ralph-C Bayer and Changxia Ke



Rockets and Feathers in the Laboratory

Ralph-C Bayer* & Changxia Ke†

April 2010

Abstract

Consumers often complain that retail prices respond faster to increases in wholesale prices than to decreases. Despite many empirical studies confirming this “Rockets-and-Feathers” phenomenon for different industries, the mechanism driving it is not well understood. In this paper, we show that, in contrast to the theoretical prediction, asymmetric price adjustment to cost shocks, as well as price dispersion, arises in experimental Diamond (1971) markets. The analysis of individual behavior suggests that the observed price dispersion can be explained by bounded rationality, as our data are well explained by Quantal Response Equilibrium (McKelvey and Palfrey 1995). Asymmetric price adjustment is caused by the updating speed of buyers with adaptive expectations being different after positive and negative cost shocks.

Keywords: Asymmetric Price Adjustment, Search Cost, Price Dispersion, Bounded Rationality, Quantal Response Equilibrium

JEL codes: D82, D83, C91, L13

*School of Economics, University of Adelaide, SA 5005, Australia.

†Max Planck Institute for Intellectual Property, Competition and Tax Law, Department of Public Economics, Marstallplatz 1, 80539 Munich.

1 Introduction

Consumers usually complain that the retail gasoline price responds faster to increases in wholesale prices than to decreases, especially when the market is volatile. However, whether this observation is a matter of fact or just a biased perception calls for empirical tests. Karrenbrock (1991), Duffy-Deno (1996) and Borenstein et al. (1997) all study the US gasoline market and conclude that their data provide strong evidence that this perceived phenomenon exists. Moreover, Borenstein et al. (1997) find that this asymmetry does not only occur in the adjustment process of retail prices to changes in wholesale prices, but also in the adjustment of spot oil prices to changes in crude oil prices. The evidence is similar for Canada (Eckert 2002) and for some European countries (Bacon 1991; Galeotti et al. 2003).

The gasoline industry is not the only industry where asymmetric price adjustment to cost changes occurs. Hannan and Berger (1991) and Neumark and Sharpe (1992) find that banks adjust both mortgage rates and consumer deposit rates asymmetrically when the central bank changes its interest rate. The adjustment typically exhibits upward flexibility and downward rigidity on mortgage rates, whereas the opposite is true for deposit rates. Therefore, the banks take advantage of both directions of interest movements by the central bank. In both the gasoline and the banking industry, input-price movement is observable to (alert) consumers, which explains why consumers recognize the asymmetric nature of the adjustment. However, for some other industries that impact even more on consumers' daily life (e.g., meat and vegetables), consumers can hardly observe input price fluctuations and hence the asymmetry is not easily observed. Ward (1982) and Goodwin and Harper (2000) confirm that also in these industries asymmetric price adjustment is the rule. Peltzman (2000) generalizes this line of research to 77 consumer goods and 165 intermediate goods across different industries and finds the asymmetry in more than two-thirds of the markets.

Although these empirical studies uncover asymmetric price adjustment to cost shocks as a stylized fact, economic theory that can explain this phenomenon is only emerging. We will review some of these recent theories in the next Section. In

traditional microeconomic theory though, variations of input prices affect the output price through marginal cost. The transmission from marginal cost to the price is governed by market power. The direction of the cost shocks does not play a role. The discrepancy between the empirical prevalence of asymmetric price adjustment and the prediction from textbook theory requires attention. This paper reports on a series of experiments designed to shed some light on the mechanisms behind the “Rocket-and-Feathers” phenomenon. Our strategy is to implement an environment, where asymmetric price adjustment should not occur according to theory and to check if we still observe it. Our ultimate aim is to identify behavioral factors that might lead to the phenomenon. Our experiments are based on the simplest market environment (two sellers and one buyer) with costly buyer search and exogenous cost shocks. Homogeneous sellers set their prices given a common production cost. A buyer (demanding one unit of the product) observes one price for free and has to decide whether to buy or to incur some search cost to learn the other price. As in Diamond (1971), who first studied such an environment, conventional theory predicts that producers charge the monopoly price for any level of positive search cost. This result seems counter-intuitive, as in the absence of search cost competitive prices would be predicted. This extreme impact of even a tiny search cost is typically referred to as the “Diamond Paradox”.

In order to study the occurrence of asymmetric price adjustment we introduce cost shocks. Initially, the production cost is commonly known to both sellers and the buyer. After a few periods a random production cost shock may occur. With equal probability the cost increases, stays the same or decreases. Sellers learn the realization of the shock before they set their prices in the after-shock markets, while the buyers only know the stochastic process which governs the shock. The introduction of the shock does not change the prediction that monopoly prices should be charged. Since the monopoly price in our basic setting is always equal to the buyer’s valuation (which does not vary with the production cost), sellers should still charge the same price regardless of the realization of the shock. The theory predicts no adjustment in our environment.

We find that play in the laboratory deviates systematically and persistently

from the theoretical predictions. Prices are well below the monopoly price and widely dispersed over the range from production cost to monopoly price. After the shock, the prices adjust in the direction of the cost changes with very different speeds depending on the direction of the shock. Prices jump up immediately after a positive shock. In contrast, prices hardly adjust in the first period after a negative shock, but fall gradually thereafter. Despite the unequal adjustment speed, the long run adjustment in magnitude seems to be symmetric. We show that bounded rationality and adaptive expectations of buyers can explain these findings. The widespread intuition that collusion is necessary for asymmetric price adjustment is rejected, as we find asymmetric price adjustment with prices well below the equilibrium price.

The paper is organized as follows. Section 2 discusses the rationale for our experimental design, points out our contribution and relates our work to existing theoretical and experimental studies. Section 3 lays out the theoretical framework and derives theoretical predictions. Section 4 describes the experimental procedure. In Section 5 we report the results on both aggregate and individual level. In Section 6 we show that bounded rationality and adaptive expectations nicely fit the data and explain our findings. Section 7 summarizes the paper.

2 Background and findings

The first intuition that springs to mind why prices adjust asymmetrically to cost shocks is (tacit) collusion. Borenstein et al. (1997) based on Tirole (1988) argue that in a world with imperfect monitoring and multiple equilibria, firms can coordinate on a high-price equilibrium using past prices as the focal point. Then firms will stick to past prices when costs fall. If costs increase old prices might not be profitable anymore, which leads to the coordination on a new price. Recently, some theoretical work by Lewis (2004), Yang and Ye (2008) and Tappata (2009) has shown that this mechanism is not necessary for asymmetric price adjustment to occur. These papers examine markets with costly price search by consumers and in-

formation asymmetries.¹ The mechanism that generates the “Rocket-and-Feather” phenomenon is based on a) price dispersion and b) different updating speeds after positive and negative cost shocks. Yang and Ye (2008) assume a continuum of firms with a capacity constraint and a continuum of consumers with three types of search cost. Equilibrium price dispersion (on two prices) occurs in the static model. In contrast, Tappata (2009) assumes a finite number of firms and buyers (with heterogeneous search cost), which also generates price dispersion (with a finite number of prices charged). Both papers assume Markov processes with some persistence to govern the cost dynamics, in which asymmetric price adjustment occurs naturally through the updating process of heterogeneous consumers. The heterogeneity of consumers is crucial here. Alternatively, instead of assuming rationality of both sellers and buyers, Lewis (2004) developed a reference-price search model where the buyers hold adaptive expectations of the market price distributions in a dynamic model. Assuming price dispersion (without actually endogenizing it), Lewis suggests that when the cost shock is positive the sellers are forced to raise the prices immediately (as the profit margin is depressed) and consumers search more when prices are higher than expected. However, if the shock is negative, the sellers only need to reduce the price enough to prevent search. The search intensity stays at similar levels when prices are equal or slightly lower than expected, and hence updating happens at a much slower pace when costs have fallen.

Our experimental design is guided by empirical questions arising from the theoretical work reviewed above. Firstly, do we need tacit collusion in order to observe asymmetric price adjustment? For this reason we chose a setup where collusion does not play a role; in our setting theory predicts monopoly prices already. Secondly, do we need an environment where price dispersion arises from equilibrium behavior of fully rational consumers and firms? Consequently, we use an environment where price dispersion is not an equilibrium in order to see if we still see price dispersion in

¹There also exist theories developed by macroeconomists, suggesting that asymmetric adjustment occurs if firms (facing nominal shocks) have to pay menu cost to adjust the prices under an inflationary trend (Tsiddon 1993; Ball and Mankiw 1994; Ellingsen et al. 2006). However, Chen et al. (2008) find that asymmetric price adjustment still occurs in the field in periods when there is no inflation.

the laboratory. Thirdly, assuming that we observe price dispersion, do we need an environment with heterogeneity where rational expectations and updating leads to asymmetric adjustment? Alternatively, does an environment suffice where rational expectations lead to symmetric adjustment, as subjects in the laboratory have adaptive expectations and cause asymmetric adjustment through them? To answer this last question we chose an environment without any asymmetries (i.e. firms are identical and there is only one consumer per market), where rational expectations prevent asymmetric adjustment speeds. The resulting design is highly unfavorable for asymmetric price adjustment, as none of the ingredients recent theories use to explain it is present. Neither Collusion, price dispersion nor different updating speeds because of heterogeneous agents with rational expectations are predicted to occur in our setting.

The market environment we implement is such that it creates a Diamond equilibrium, which gives a crisp prediction: unitary monopoly prices. To fill the gap between this extreme and the other of marginal cost pricing in Bertrand competition, a large number of theoretical models (which generate equilibrium price dispersion) were developed (e.g., Reinganum 1979; Braverman 1980; Varian 1980; Burdett and Judd 1983; Carlson and McAfee 1983; Rob 1985; Stahl 1989). Equilibrium price dispersion arises in these models mainly due to the introduction of heterogeneity on either sellers (in production cost) and/or the buyers (in search cost or search technology). The experimental evidence on the Diamond Paradox is mixed. Grether et al. (1988) observed prices close to the monopoly price in three out of four sessions they conducted. More recent studies (David and Holt 1996; Abrams et al. 2000) find evidence that search cost increases prices but does not lead to the monopoly outcome. Interestingly, Cason and Friedman (2003) observe that prices are close to monopoly prices if the buyers are played by computer automata, while they are much lower if the buyers are people.

The knife-edge nature of the Diamond equilibrium relies on two classical assumptions typically made by solution concepts for dynamic games of imperfect information. First, the buyer has the correct belief that both sellers charge the same price in equilibrium, which rules out searching. Second, the sellers have common knowledge

of rationality, which implies that each player is able to compute best responses and also knows that the other players are. In contrast to this, analyzing individual behavior in our experiment suggests the existence of a significant portion of bounded rationality (or noisy behavior) of both sellers and buyers. Baye and Morgan (2004) show that allowing for bounded rationality leads to price dispersion in a standard Bertrand Oligopoly, as observed in experimental studies (Dufwenberg and Gneezy 2000; Abrams et al. 2000). Compared to the simple Bertrand model, the Diamond environment is more complex and an even higher level of cognitive ability is required to find the optimal strategy. Given the higher level of complexity it is also less likely that a player believes the others have adequate cognitive abilities to solve for equilibrium. We conjecture that bounded rationality and/or the lack of common knowledge of rationality is likely to be at work in the Diamond world. Our data shows that the sellers' pricing strategies (and the buyers' searching decisions) are consistent with Quantal Response Equilibrium (QRE; McKelvey and Palfrey 1995), which models bounded rationality by assuming that players' choices are stochastic where the probability of playing a strategy increases with the expected payoff.

Once we allow for noise in best-responses (like in QRE) price dispersion occurs naturally. If price dispersion is present then the buyers need to form beliefs about the price distribution in order to decide if searching is expected to be profitable. Buyers in the laboratory are found to be adaptive learners, who form their expectations according to the prices observed in previous markets. After a cost shock when production cost becomes the sellers' private information, more searching speeds up the buyers' updating process of the true cost state. After a positive shock it is in the interest of the sellers to convince buyers that the cost has gone up as quickly as possible, as they want to make sure that buyers quickly learn that prices now will be higher in general. As a consequence prices jump up immediately. Facing a market-wide price jump, adaptive buyers initially search, as they believe that the second price is likely to be lower. The high frequency of search leads to rapid updating. After a negative shock it is in the interest of the sellers to slow down the diffusion of knowledge on the direction of the cost shock. Sellers initially keep their prices at the pre-shock level. Buyers with adaptive expectations have no strong incentive to

search if the prices stay where they were before. Updating is slow as there is not much searching. However, gradually the information that the cost has fallen filters through. Firms adjust their prices accordingly. Price adjustment is asymmetric. In summary we find that a) tacit collusion is not necessary for asymmetric price adjustment, b) the market does not have to be such that (rational) theory predicts price dispersion for asymmetric price adjustment to be observed, and c) no specific kind of heterogeneity is needed to generate different updating speeds after positive and negative cost shocks.

3 Theoretical framework

We adopt a simple two-phase dynamic market game, which incorporates search cost, information asymmetries and exogenous cost shocks. We will refer to the two phases as *pre-shock* and *after-shock* phase. Without loss of generality we assume that each phase consists of a finite number of k markets. In each market, there are two sellers and one buyer. Each seller intends to sell one unit of a homogeneous product which costs mc to produce. The buyer demands only one unit of this product and values it at $v > mc$.

The timing in the market is the following. First, sellers – indexed 1 and 2 – independently and simultaneously set prices p_1 and p_2 . After both sellers have set their prices, one of these prices is randomly drawn and displayed to the buyer for free. Having observed the free sample price p_i the buyer can either *buy* at that price, *search*, or *exit* the market. The market ends immediately, if the buyer choose to *buy* or *exit*. However, if *search* is chosen, the buyer learns the other firm's price p_{-i} . Searching is costly though. A search cost c is incurred by a buyer who chooses to *search*. Having observed both prices, the buyer can choose to buy at any of the two prices, or exit the market.

The buyers' valuation, search cost and production cost are common knowledge and remain constant in all *pre-shock* markets. Between the *pre-shock* and *after-shock* phase an industry-wide cost shock may occur, i.e., the production cost mc may take on a different value. The realization of the post-shock marginal cost is

privately observed by all sellers, whereas buyers only know the stochastic process which governs the shock.² Apart from the cost shock and the arising private knowledge of the sellers, the *after-shock* markets are otherwise the same as the *pre-shock* markets.

A seller's payoff in a particular market is $\pi_i = p_i - mc$ given that he can sell the good. Otherwise, π_i is equal to 0. The buyer's payoff (conditional on buying) in each market is equal to her valuation, less the price paid and the search cost (if search takes place). In situations where the buyer exits the market, the payoff is either 0 or $-c$ depending on whether *exit* occurs before or after *search*.

3.1 The monopoly equilibrium

It is obvious that our model yields a perfectly competitive Bertrand equilibrium if search is free (i.e., $c = 0$). However, the equilibrium shifts to the other extreme (the monopoly equilibrium) if there exists only a small positive search cost. In order to prove this we need to introduce some notation. Denote the price seller i sets in market $t \in \{1, 2, \dots, T\}$ as p_{it} . Further denote the density function of the stochastic process that determines the marginal cost in period t as $f_t(mc)$.³ The price a buyer sees first (before searching) is denoted by p_t^1 , while the second price (after search) is given by p_t^2 . The choice a buyer makes after observing the first price is $a_t^1(p_t^1)$ and the choice after searching is $a_t^2(p_t^1, p_t^2)$. The beliefs the buyer holds about the price he will discover after *search* (conditional p_t^1) is written as $\mu(p_t^2 | p_t^1)$, where μ is a density that is defined for all possible prices.

Proposition 1 *For any $c > 0$ and any stochastic process which governs mc (with p.d.f $f_t(mc)$ with support $[\underline{mc}, \overline{mc}]$, $\overline{mc} \leq v$), there exists a unique pure-strategy*

²It is common knowledge that the shock is industry wide, i.e. all firms have the same production cost after the shock.

³Note that this allows for an environment richer than what was implemented in the experiments with just one shock period. Additionally, we could allow the stochastic process to condition on previous realizations of mc , which would not change our result.

Perfect Bayesian Nash Equilibrium that also satisfies subgame perfection, where

$$\begin{aligned}
p_{it}^* &= v \quad \forall i \in \{1, 2\}, t \in \{1, 2 \dots T\}, \\
a_t^{1*}(p_t^1) &= \begin{cases} \text{exit} & \text{if } p_t^1 > v \\ \text{buy} & \text{if } p_t^1 \leq v \end{cases} \quad \forall t \in \{1, 2 \dots T\}, \\
a_t^{2*}(p_t^1, p_t^2) &= \begin{cases} \text{exit} & \text{if } \min[p_t^1, p_t^2] > v \\ \text{buy at } p_t^1 & \text{if } p_t^1 \leq \min[v, p_t^2] \\ \text{buy at } p_t^2 & \text{if } p_t^2 \leq \min[v, p_t^1] \end{cases} \quad \forall t \in \{1, 2 \dots T\}.
\end{aligned}$$

give the strategies.

Proof. Following backward induction in order to ensure subgame perfection, we first solve the last market T . We now investigate if there exists an equilibrium in which seller 1 charges p_{1T}^* and seller 2 charges p_{2T}^* with $p_{1T}^* \geq p_{2T}^*$ ⁴ by looking at three mutually exclusive cases covering all possible prices. In what follows we drop the subscript T for ease of notation.

(i) Suppose there existed an equilibrium with $p_1^* - p_2^* > c$, then on the equilibrium path

$$a^{1*}(p^1) = \begin{cases} \text{search} & \text{if } p^1 = p_1^* \\ \text{buy} & \text{if } p^1 = p_2^* \end{cases}$$

and

$$a^{2*}(p^1, p^2) = \begin{cases} \text{exit} & \text{if } \min[p^1, p^2] > v \\ \text{buy at } p^1 & \text{if } p^1 \leq \min[v, p^2] \\ \text{buy at } p^2 & \text{if } p^2 \leq \min[v, p^1] \end{cases} \quad (1)$$

is the optimal strategy for the buyer with the correct beliefs about the second price once he has seen the first

$$\mu_t(p_1^* | p^1 = p_2^*) = 1 \quad (2)$$

$$\mu_t(p_2^* | p^1 = p_1^*) = 1. \quad (3)$$

⁴The identity of the seller charging the higher price is not important.

Given this optimal response of the buyers, the *ex ante* expected equilibrium profit of seller i is given by

$$E\pi_i(p_i^*, p_{-i}^*, a^{1*}, a^{2*}) = \begin{cases} 0 & \text{if } i = 1 \\ (p_i^* - mc) & \text{if } i = 2 \end{cases}.$$

Observe that seller 1 always has an incentive to deviate by mimicking seller 2 and charging p_2^* . His expected profit then would be equal to $(p_2^* - mc)/2$. There is no stage game equilibrium with $p_1^* - p_2^* > c$.

(ii) Suppose there exists an equilibrium with $c \geq p_1^* - p_2^* > 0$, then the sequentially rational continuation on the equilibrium path is

$$a^{1*}(p^1) = \text{buy if } p^1 \in \{p_1^*, p_2^*\}^5$$

This time seller 2 has an incentive to mimic seller 1 as increasing the price to p_1^* increases the expected payoff from $(p_2^* - mc)/2$ to $(p_1^* - mc)/2$. So $p_1^* - p_2^* \leq c$ cannot hold in any equilibrium.

(iii) From the above we know that $p_1^* - p_2^* = 0$ (i.e., $p_1^* = p_2^* = p^*$) must hold in any equilibrium. We can rule out p^* below mc or above v , as sellers have non-positive expected payoffs. Sellers can always deviate either upwards or downwards to get a positive expected payoff. Therefore, we can concentrate on $p_i^* \in [mc, v]$. Now suppose that there exists an equilibrium with $p_1^* = p_2^* < v$. Then the sequentially rational continuation for the buyer is

$$a^{1*}(p^1) = \begin{cases} \text{buy} & \text{if } p^1 \leq p^* + c \\ \text{search} & \text{if } p^1 > p^* + c \end{cases}$$

with beliefs

$$\mu_t(p^* | p^1) = 1$$

Then for any seller deviating to a price $p_i = p^* + \varepsilon$ increases the expected profit from $(p^* - mc)/2$ to $(p^* + \varepsilon - mc)/2$ as long as $p^* < v$. While at $p_1^* = p_2^* = v$, any

⁵The continuation off the equilibrium path after searching is given by (1). The correct beliefs are given by (2) and (3).

deviation will reduce the expected profit. Consequently, we have $p_1^* = p_2^* = v$ as the unique equilibrium in the market.

Now observe that the equilibrium is unique with respect to the equilibrium payoffs and independent from histories and realizations of the beliefs about mc . Therefore, in equilibrium this stage game equilibrium will be played in all market stages.

■

Such a monopoly equilibrium exists not only in our simple model but also in the more general search models with more sellers and buyers. The basic intuition goes back to Diamond (1971). Introducing an infinitesimally small search cost into an otherwise perfectly competitive market can make firms become local monopolists. The optimal search strategy is that a rational buyer should only search if his expected gain from search is greater than the search cost. Keeping the buyer's strategy in mind, a seller always has an incentive to charge a higher price than the other firm as long as the deviation does not induce consumer search. For any price lower than the monopoly price, sellers have an incentive to raise the price, where the rise has to be smaller than c and the new price has to be weakly smaller than the customers' valuation. This deviation process continues until the monopoly price is reached. In equilibrium a uniform monopoly price prevails and the sellers neither have the incentive to raise nor to lower the price. Given our specific setup, the equilibrium price is at v , regardless of production cost and buyers' beliefs about them. This extremely simple model yields a sharp prediction, which greatly facilitates the comparison between the laboratory results and the theoretical prediction. It also provides a great environment for identifying the behavioral factors that play a role in asymmetric price adjustment, since the model predicts no adjustment at all.

4 The experiment

Our experimental design and procedure follows the theoretical framework described above closely. We set $v = 100$, $c = 15$ in the experiment. The marginal cost mc is set to 30 in the *pre-shock* phase. In the *after-shock* phase, v and c remain the same, while mc might change. Due to some exogenous shock, the production cost

will *increase from 30 to 50, decrease from 30 to 10, or may not be affected at all*. Each of these three events is equally likely to occur. Therefore, the production cost in the *after-shock* markets take on any one of the three values (50, 10, 30) with probability one-third. Parameters and probabilities are made public in advance. After the shock, the new cost value is announced to the sellers only. Given the realization of the cost shock we have three different treatments which we refer to as *MC-Increase*, *MC-Decrease* and *MC-Constant*, respectively. Despite the variation in *mc* and the resulting private information in the *after-shock* markets, the equilibrium prices predicted by theory (assuming fully rational and selfish agents) are constant at 100 in all treatments. The parameter values are summarized in Table 1.

Treatment	MC-Increase	MC-Decrease	MC-Constant
Valuation (v)	100	100	100
Search cost (c)	15	15	15
Pre-shock Cost (mc_{1-15})	30	30	30
After-shock Cost (mc_{16-30})	50	10	30
Predicted equilibrium price	$p^* = 100$	$p^* = 100$	$p^* = 100$

Table 1: A summary of parameter values by treatment.

The experiments were conducted at AdLab at the University of Adelaide. Subjects were recruited from university students in various disciplines and at various stages of their tertiary education. The experiments were programmed and conducted using Z-tree (Fischbacher 2007). In total, 168 subjects participated in 10 different sessions with no repeated participation.

In each session, one specific treatment was randomly assigned. Subjects were given written instructions, which they had time to read before the experiments commenced. Subjects were randomly assigned their roles (seller or buyer) at the beginning of the experiment. Players' roles remained the same throughout the whole session. Each session consisted of 30 markets with 15 markets each in the *pre-shock* and *after-shock* phase. We chose typed-stranger matching in order to eliminate repeated-game effects. Consequently, subjects were newly and randomly matched in a group of three (two sellers and one buyer) in each new market.

In each market, sellers had only one decision to make (i.e., set the price). After all sellers had set their prices, the buyers entered the market. The buyers each had

one or two decisions to make depending on whether they decided to search or not. At the end of each market, profits were calculated and displayed to the subjects before a new market (after random re-matching) started. Sellers did not know the prices posted by other sellers in previous markets. However, information on whether the buyer had seen their price was provided in the profit-display stage. In between the two phases (i.e., between period 15-16), subjects were reminded that a cost shock might have occurred. Subjects were reminded of the potential states of the world (different marginal cost levels), the probability distribution over the states and the fact that the realized state would only be learned by the sellers. On average, each session took around 1 hour and 15 minutes, during which subjects earned about 19 Australian Dollars on average.

5 Results

While standard theory predicts a unitary monopoly price (at $v = 100$) in both *pre-shock* markets and *after-shock* markets, the experimental data paint a completely different picture. The two regularities observed in the data are: (i) the prices are very dispersed below the monopoly price; and (ii) although the prices adjust in the direction of the cost movement, the adjustment is much faster after a positive cost shock than after a negative shock.

In this section we establish these stylized facts by presenting the data on an aggregate level. A deeper analysis on the individual level follows, which aims at identifying the driving factors behind the deviations.

5.1 Price dynamics

Behavior in the laboratory closely resembles the aggregate price dynamics identified by empirical studies in real-world industries exhibiting the “Rocket-and-Feather” phenomenon. This can be seen in Figure 1, which depicts the average posted-price time series over 30 periods by treatment. In the *pre-shock* phase the three treatments essentially produce identical price time series. Not surprisingly, we observe similar price patterns in all treatments before the shock. The average posted prices all start

at about 60 to 65 in period 1 and then slowly trend down to 55 to 60 in period 15.

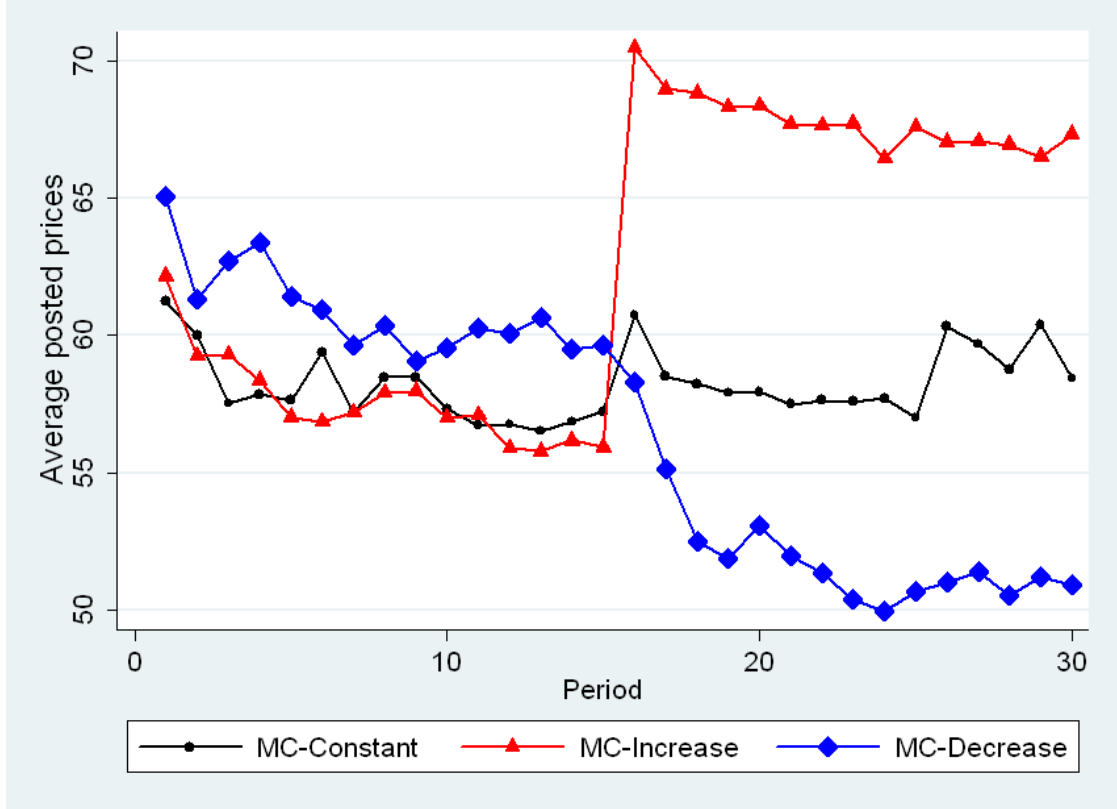


Figure 1: Average posted prices by period and treatment.

After the shock, prices adjust in the direction of the cost changes. This is in contrast to the theoretical prediction that actual prices should be independent of production cost. It is notable that the adjustment magnitude and speed in the positive shock treatment considerably differ from those in the negative shock treatment. The adjustment is instantaneous following a positive shock, while the same magnitude of adjustment takes at least seven periods after a negative shock. When the production cost remained unaltered, the price pattern in the *after-shock* phase is similar to that in the *pre-shock* phase.

The different adjustment speeds after positive and negative shocks can be seen clearly in Figure 2. Taking the average posted prices in period 15 (p_{15}) as a benchmark, we plot the absolute difference between the mean posted prices in each period after the shock and the mean posted price in period 15 (i.e., $|p_t - p_{15}|$).⁶ We see that

⁶We plot absolute deviations, as we are interested in the adjustment speed rather than direction here.

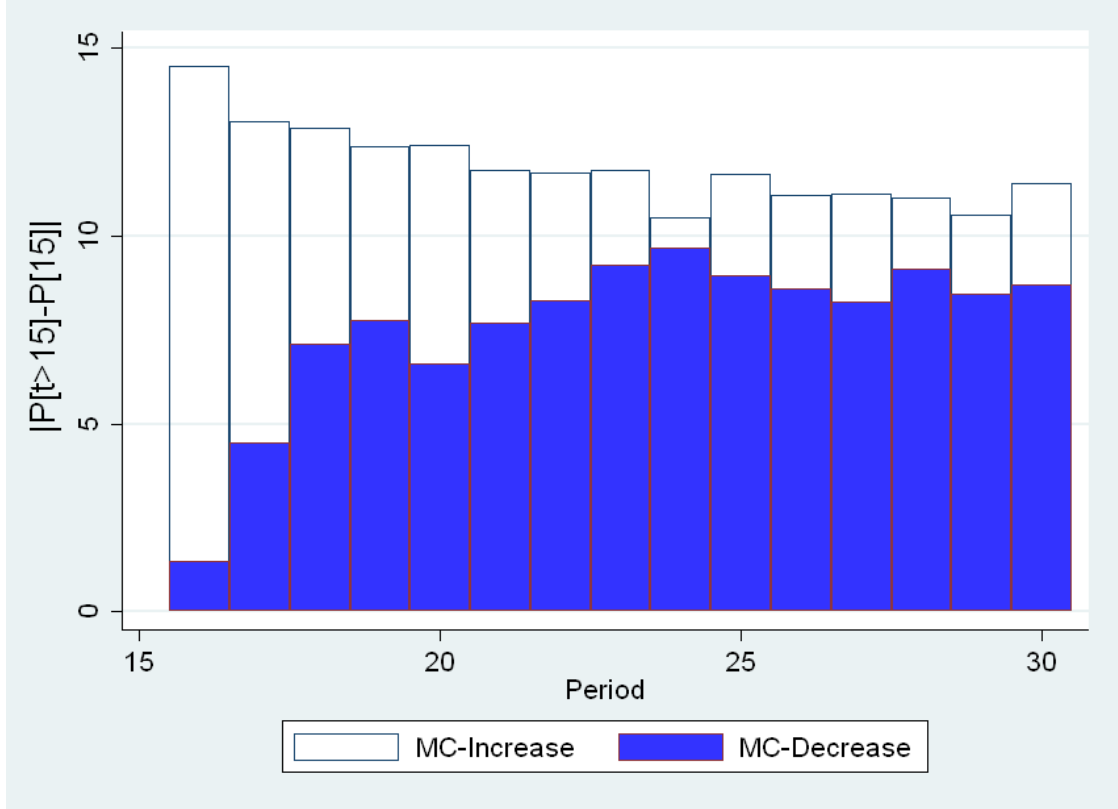


Figure 2: Price deviations from the pre-shock level.

average posted prices jump up immediately by around 14.5 units in the *MC-Increase* treatment, whereas the immediate average adjustment is only 1.3 units after a negative shock. Note that the magnitude of the cost shock is the same (20 units) in both treatments. The short-run adjustment overshoots when the production cost increases, as the average price deviation (from pre-shock levels) decreases after the initial shock period. In contrast, the short-run adjustment after a decrease in cost is sluggish, but prices generally decrease over time (until nine periods after the shock). Despite the initially large gap in the adjustment speed, the size of the adjustment for positive and negative shocks tends to converge in the long run. In the long run about half of the cost changes are reflected in the price, which means that the change in surplus is shared equally by sellers and buyers. These findings are further supported by non-parametric tests on individual-level data. A Wilcoxon rank-sum test (two sided) rejects the null hypothesis that the adjustment (in absolute values) is equal in the *MC-Increase* and *MC-Decrease* treatments from period 16 to 22 on the five-percent level, but not afterwards.

One interesting question is how firms adjust their prices depending on what happened the period before. Note that a seller could not learn too much from the history, as she only knows if she sold the good, did not sell the good after the customer had seen her price, or if she did not sell and the customer did not visit her shop. Furthermore, our typed-stranger matching limits the value of what the firm can infer from this observation as the market participants in the next round will be different. The median price changes between periods for the cases “sold” and “unsold-unseen” is plus one (sold in the MC-Increase treatment) or zero (all other cases and treatments). The mean price changes for “sold” are slightly positive in all treatments though (, as very few firms reduced their prices after selling). The only situation where a firm can clearly learn anything about how her price compares is when the customer visited the shop but then bought from the other seller. In this case the median reaction we observe is a price reduction by five monetary units. This is the same in all three treatments. This shows how important search is for competitive dynamics, which tells us that search behavior after a cost shock is key to explain the price dynamics. Only if firms who did not sell know that the consumer has seen both prices a sizable adjustment follows. As this can only happen if the consumer searched, over-all search intensity is important for downward movements of prices.

5.2 Price distributions and search

In order to understand the deviations from the theory predictions described above, we now investigate the data on a less aggregated level. Sellers set prices without knowing whether they are going to be seen first or second (i.e., p_1 or p_2). Therefore, p_1 is a random sub-sample of all posted prices: it is a good representation of all prices set by sellers. In what follows we plot the distribution of these free prices by treatment and period and also identify the prices which triggered search. This provides some insight into individual pricing and search behavior.

Figures 3, 4 and 5 show scatter plots of the free observed prices (p_1) by treatment over time. The plots show that sellers never price below their production cost (mc)

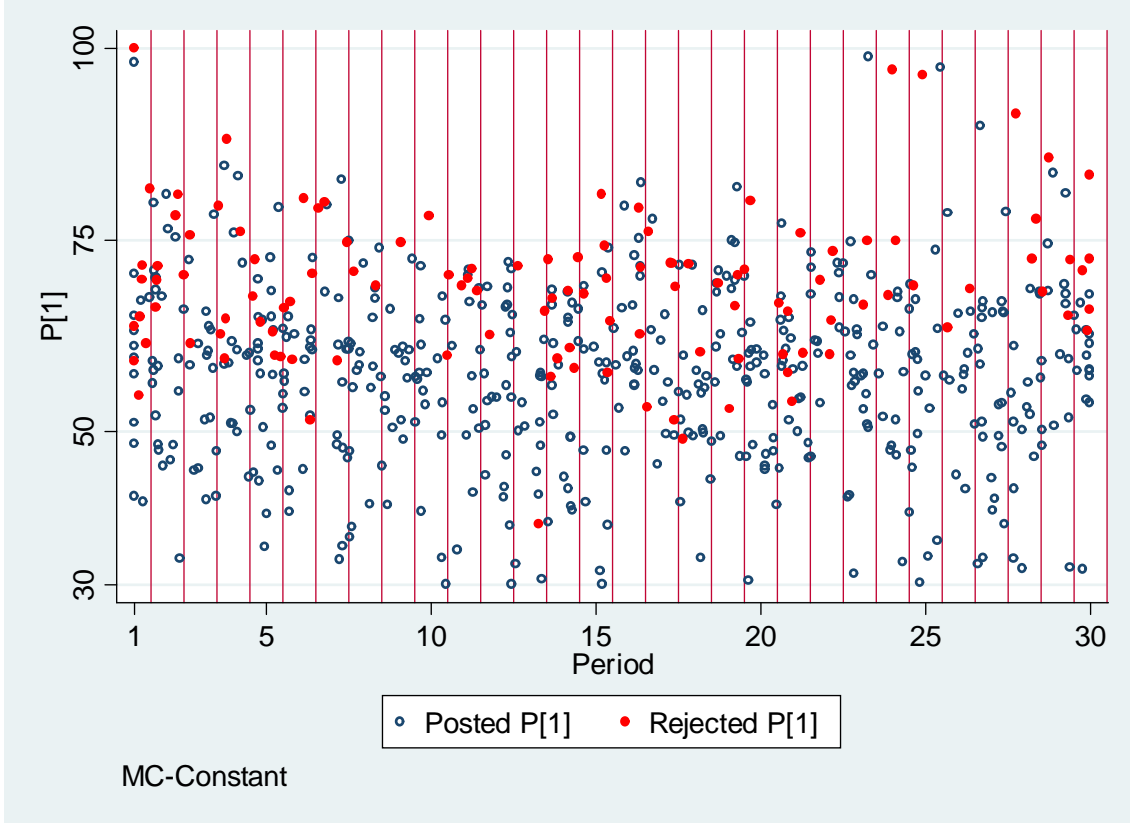


Figure 3: Posted prices and search in the *MC-Constant* treatment

and hardly ever price above the buyers' valuation.⁷ Instead of all sellers posting the same monopoly price as predicted by theory, we observe price dispersion within the range of $[mc, v]$ in all periods, without any significant tendency to converge to a single price. In the *pre-shock* phases, the price distributions share the same characteristics and dynamic patterns in all treatments. The only notable difference is that the prices in the *MC-Decrease* treatment are slightly less dispersed. However, robust equal-variance tests do not reject the null hypothesis of equal variances in 13 of the 15 periods.

After a positive shock, prices are naturally bounded from below by the new production cost (50). We therefore observe a clear upwards shift of the whole distribution, followed by buyers searching in more than 80% of the cases in the market immediately after the shock. The search intensity drops sharply in the next period

⁷Actually, only in four markets (involving two particular subjects out of 168) were offers higher than 100 posted. We exclude these four outliers in the figures above in order to maintain a reasonable scale.




Figure 4: Posted prices and search in the *MC-Increase* treatment

for similar prices. This suggests that customers have concluded from their initial search that all prices have risen. Thereafter search patterns quickly stabilize at a low intensity. Consequently, there is no need for the sellers to dramatically revise their prices downwards in later periods with the effect that prices also settle quickly after the shock. In stark contrast, we do not observe this dramatic instantaneous shift from the *pre-shock* phase to the *after-shock* phase, when the underlying production cost decreases. The price distribution gradually moves downwards and becomes more dispersed over time.

5.3 Buyers' search rule

Having analyzed the sellers' pricing strategy, we now turn to a more detailed analysis of buyer behavior. In contrast to the theoretical prediction (zero search in equilibrium), buyers search in about a quarter of the markets. Note that given that prices are dispersed searching at high first prices becomes sequentially rational. *Exit* is

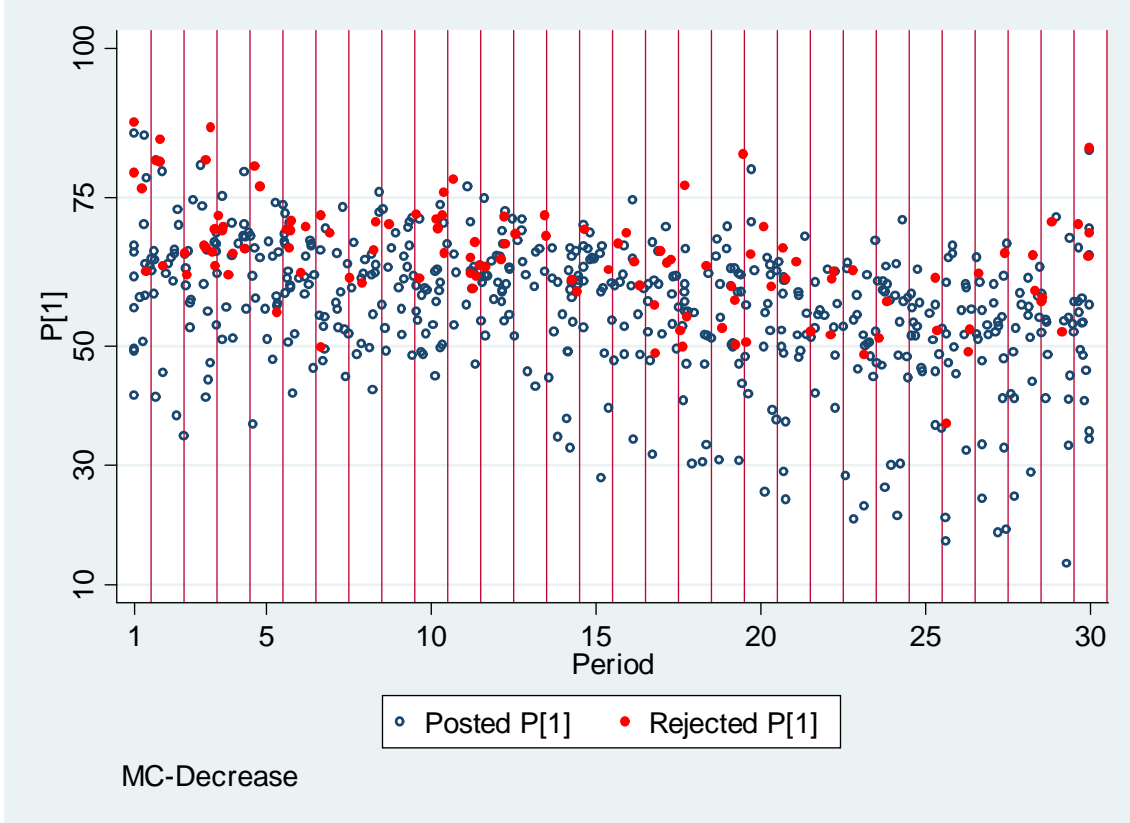


Figure 5: Posted prices and search in the *MC-Decrease* treatment

not sequentially rational for any price below 100 though. We only observe irrational *exit* in 19 out of the 1680 cases.⁸

If our buyers all had identical risk preferences and were fully rational and held correct beliefs on the price distributions then we should observe a clean cutoff rule. If the price is above a certain cutoff price for given beliefs everyone searches and for prices below no one does. Figures 3, 4 and 5 showed that this is not the case. We use a random-effect logistic panel model in order to estimate buyers' search rules, allow for some heterogeneity and noise and to identify trends and cost-shock effects. With the dependent variable being the probability of search our covariates of particular interest are:

- P_1 - the initially price observed

⁸The rare exit decisions help us to rule out inequity-aversion as an explanation for the prices below v . Camerer (2003) reports on a large number of experimental ultimatum games. A consistent finding is that offers below 30% are rejected about 25% of the time. Our data, where prices, which leave the consumer with less than 30% of the surplus, are only rejected in 3% of the cases, rule out inequity aversion driving our results.

- *Aftershock* - a dummy variable to separate the *after-shock* phase observations from the base category *pre-shock* observations.
- P_1* *aftershock* - interaction term of P_1 and *Aftershock*.
- T_t^+ - a set of dummy variables for each period after 16 (i.e., $t > 16$) in the positive shock treatment.
- T_t^- - a set of dummy variables for each period after 16 (i.e., $t > 16$) in the negative shock treatment.

Variable	β	Std.Err.	$\frac{\partial P}{\partial X}$	\bar{x}	Variable	β	Std.Err.	$\frac{\partial P}{\partial X}$	\bar{x}
α	-14.974***	1.322			P_1	0.235***	0.0193	0.027***	59
Shock	1.285	1.640	0.091	1	P_1* Shock	-0.021	0.026	-0.002	59
T_{17}^+	-1.052*	0.606	-0.082**	0	T_{17}^-	1.421***	0.671	0.255*	0
T_{18}^+	-0.887	0.617	-0.073*	0	T_{18}^-	1.062**	0.754	0.173	0
T_{19}^+	-1.977***	0.660	-0.111***	0	T_{19}^-	-0.276	0.967	-0.029	0
T_{20}^+	-1.270**	0.594	-0.091***	0	T_{20}^-	1.986***	0.758	0.394**	0
T_{21}^+	-1.511**	0.620	-0.100***	0	T_{21}^-	0.613	0.800	0.087	0
T_{22}^+	-1.919***	0.662	-0.110***	0	T_{22}^-	1.195	0.757	0.202	0
T_{23}^+	-2.415***	0.745	-0.119***	0	T_{23}^-	1.323*	0.761	0.231	0
T_{24}^+	-1.524**	0.641	-0.100***	0	T_{24}^-	-0.321	1.145	-0.033	0
T_{25}^+	-2.634***	0.774	-0.121***	0	T_{25}^-	0.355	0.969	0.046	0
T_{26}^+	-2.351***	0.717	-0.118***	0	T_{26}^-	0.105	0.928	0.013	0
T_{27}^+	-3.317***	0.867	-0.127***	0	T_{27}^-	1.235	0.851	0.211	0
T_{28}^+	-2.557***	0.750	-0.120***	0	T_{28}^-	0.117	0.983	0.013	0
T_{29}^+	-2.608***	0.801	-0.121***	0	T_{29}^-	0.783	0.887	0.118	0
T_{30}^+	-2.508***	0.763	-0.120***	0	T_{30}^-	1.342*	0.773	0.236	0

Log likelihood = -622.91747

Wald chi2(41) = 260.03***

rho 0.243***

Note: Other control variables (course, age, maths background) are not significant. The marginal effect ($\frac{\partial P}{\partial X}$) is predicted at the specified \bar{x} and mean values of other control variables. * p -value<0.1; ** p -value<0.05; *** p -value<0.01

Table 2: A random-effect logistic estimation of the probability to search.

The estimated coefficients and the corresponding marginal effects are presented in Table 2.⁹ The results indicate that buyers seem to search according to a probabilistic rule that has the following properties:

⁹The values of the independent variable at which the marginal effects are estimated are given in the table in column “ X ”. Some control variables (such as demographics), which we do not list here are set to their mean.

1. Buyers' probability to search increases with the initial price (p_1). There is no clear universal cutoff price leading to a sharp increase of search probability.
2. Buyers follow the same probabilistic search rule in different treatments and periods before the shock, while after the shock the search rules differ across treatments.
3. After a positive shock, buyers' probabilistic search rule does not change in period 16. Facing a much higher p_1 than in previous periods, almost all buyers choose to search in period 16 and follow their original rule. In period 17 the search rule of the subjects in the *MC-Increase* treatments changes. At the same prices buyers' search probability decreases by 0.08. This shift is persistent and can be seen by inspecting the marginal effects for the dummies T_{17}^+ through to T_{30}^+ , which are all around -0.1 . This means that the subjects' beliefs have been updated within one round of high intensity search. Therefore, after the initial search buyers believe that the price distribution has shifted upwards due to a positive cost shock and permanently adjust the search rule in period after period 16.
4. After a negative shock, with the exceptions of periods 17 and 20 the search probabilities at given prices are not significantly different (on the 5% level) from those in the *MC-Constant* treatment.
5. The pure knowledge of a shock occurring does not change the relationship between prices and search probability, as the coefficient on P_1 **Aftershock* not significant.

These results together with the observed prices are clear evidence for adaptive searching. Consumers use their observations in previous rounds in order to form expectations about the price distributions. The observation of a high prices after a positive cost shock leads to lots of searching for the given search rule. Observing that actually the other price is also high once search has happened, lets consumers adapt their search rule in the *MC-Increase* treatment instantly. In the case of the *MC-Decrease* treatment, sellers slightly lower their prices once the shock occurred.

Consumers hardly search (given their adaptive search rule), which makes it hard to update that now prices should be lower. The search rule does not properly adjust. In the *MC-Constant* treatment unchanged prices in the period after the shock and the adaptive search rule leads to business as usual for the subsequent markets. The asymmetric information injected into the market through the potential shock does not unsettle the market at all.

6 Is bounded rationality the missing link?

In the previous Section we have established some important results necessary for explaining asymmetric price adjustment. We have identified deviations in behavior from theory (like price dispersion and adaptive search rules) that are likely to parts of the explanation why we observe asymmetric price adjustments. In this Section we want to see if we can tell a consistent story. For this purpose, we try to explain our data by using a noisy equilibrium concept capturing bounded rationality. In what follows we will use Quantal Response Equilibrium (QRE) in order to test if we can explain the *pre-shock* data (price distributions and search behavior). Later we will then use the same model to capture behavior in the *post-shock* phase, once behavior has settled after the shock. In the treatments *MC-Constant* we have conjectured that no real adjustment happens, whereas for the *MC-Increase* treatment we claimed that adjustment takes place entirely within the period after the shock. So we refit the model for these two treatments *post-shock* phases without the shock period and compare the parameters measuring the level of irrationality to those from the QRE for pre-shock phase. In the *MC-Decrease* treatment we argued that the adjustment process takes longer. Therefore, we fit the model to the remaining periods and again compare parameters. A comparison of the parameter values necessary for a good fit across all these situations then allows us to conclude if one model (here QRE) can equally well explain behavior in these different phases.

6.1 Explaining price dispersion with Quantal Response Equilibrium

Throughout the experiment we observe price dispersion and average prices well below the monopoly price. Following Baye and Morgan (2004), who showed that price dispersion in Bertrand oligopolies can be explained by bounded rationality, we use Quantal Response Equilibrium to explain our data (McKelvey and Palfrey 1998). In a Quantal Response Equilibrium (QRE) subjects play noisy best responses to each other. A player plays a mixed strategy, where the action with the highest density is the best response to the mixed strategies of all the other player. However, a player makes mistakes and also plays actions with positive probability, which are not best responses. The likelihood with which a certain action is played is determined by an error structure. Typically one assumes that the larger the payoff consequences of deviating from a best response the lower is the probability for such an action to be played. The two specifications dominating the literature are the extreme-value distribution of errors leading to an exponential choice function and the power-law rule introduced by Luce (1959). We will use the latter.

First, for simplicity we assume that a consumer, once she knows both prices will choose the lower price with certainty, as it only happened in 0.5 percent of the cases that a fully informed consumer bought from the more expensive seller. Now suppose a firm faces a search rule of the consumer that assigns a search probability to any price p observed first.

$$prob\{search|p\} = g_b(p)$$

Furthermore denote the mixed strategy of the other firm (i.e. the distribution of prices chosen) as $\sigma_{-i}(p)$ with

$$\sum_{p=mc}^v \sigma_{-i}(p) = 1,$$

$$\sigma_{-i}(p) \in [0, 1] \quad \forall p \in [mc, v].$$

The restriction on $p \in [mc, v]$ assumes that firms never post a price that leads to a loss or that will never be accepted, as it would lead to a negative payoff of the

buyer. Then the expected payoff of the firm is given by:

$$E\pi_i(p, \sigma_{-i}, g_b) := (p - mc) \left[\frac{(1 - g_b(p))}{2} + \frac{g_b(p)}{2} \sum_{r=p}^v \sigma_{-i}(r) + \sum_{r=p}^v \frac{g_b(r) \sigma_{-i}(r)}{2} \right],$$

where the first part in brackets is the probability of firm i having the own price displayed first and being accepted without search. The second part is the probability that firm i 's price is shown first, the consumer searches but then accepts the price of firm i , as it is lower than that of firm $-i$. The last part is the probability that the other firm's price is displayed first, the consumer searches and then accepts the price of firm i . With a power-law rule firm i 's price distribution becomes

$$\sigma_i(p) := \frac{E\pi_i(p, \sigma_{-i})^{1/\mu_i}}{\sum_{r=mc}^v E\pi_i(r, \sigma_{-i})^{1/\mu_i}}. \quad (4)$$

The parameter μ_i captures the level of rationality. The lower μ_i the less noisy is the decision of seller i . If $\mu_i \rightarrow 0$ then the seller always picks the best response, while $\mu_i \rightarrow \infty$ leads to the seller randomizing over all prices with equal probability. In a symmetric QRE we have identical price distributions for the two identical firms and identical noise parameters, i.e.

$$\sigma_i = \sigma_{-i} = \sigma \quad (5)$$

$$\mu_i = \mu_{-i} = \mu_s. \quad (6)$$

Now turn to the consumer. Her profit from buying right away at the first shown price is

$$U_b(\text{buy}, p) = v - p,$$

while the expected profit from searching is given by

$$EU_b(\text{search}, p, \sigma) = v - p \sum_{r=p}^v \sigma(r) - \sum_{r=mc}^{p-1} r \sigma(r) - c.$$

In order to be able to use a power-law rule as above we have to transform the payoffs, as the expected profit from searching at a certain price is negative for some p and

σ .¹⁰ We add c to both payoffs, as this is the maximum loss possible:

$$g_b(p) := \frac{(EU_b(\text{search}, p, \sigma) + c)^{1/\mu_b}}{(EU_b(\text{search}, p, \sigma) + c)^{1/\mu_b} + (U_b(\text{buy}, p) + c)^{1/\mu_b}}. \quad (7)$$

Solving equations (4), (5), (6) and (7) simultaneously gives the desired QRE. Ideally, one would solve for the equilibrium and then use the data to estimate the irrationality parameters μ_s and μ_b . However, this is not feasible, as a closed-form solution does not exist. This same problem also prevents us from allowing for and estimating risk-aversion, which would be desirable (as e.g. in Goeree et al. 2002). We use numerical methods to solve for the QRE for given irrationality parameters and then compare the resulting price distribution and search rule with their empirical counterparts.¹¹ Then we repeat this exercise with different irrationality parameters in order to decide, which fits best. As the criterion we use a χ^2 statistic for the distance between empirical and QRE price distribution.¹² For the search rule we define a distance measure d , which is the weighted average absolute distances between empirical and predicted probabilities. The average is taken over all prices observed first and the weights are the frequencies with which the prices were actually charged.

6.2 Pre-shock phase

In what follows we show that the behavior before the cost shock can nicely be explained by QRE. Figure 6 shows the price distribution and search rule with a fitted QRE model and compares it to the empirical values. For the plot of the empirical price-distribution we smoothed the distribution using a kernel. The irrationality parameters that fit best are $\mu_s = 0.5$ and $\mu_b = 0.24$.¹³ As expected the level of noise is higher on the sellers' side, as the decision to set a price is more complex than the buyer's task of deciding to search or not.

¹⁰One such situation is the equilibrium prediction where all prices are equal to v .

¹¹The Mathematica code performing the iterative procedure necessary can be obtained on request from the authors.

¹²For the calculation we group prices into five bins in order to avoid cells with low counts. This is done for all treatments in order to allow for comparisons.

¹³The distance measures are $\chi^2 = 34.64$ and $d = 0.094$.

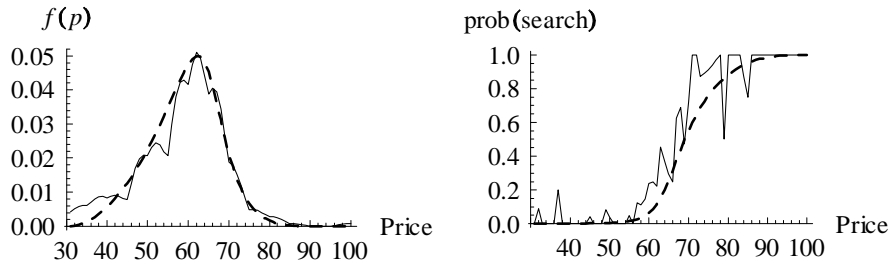


Figure 6: QRE predictions and empirical data for the pre-shock phase

Over all, we can conclude that the Diamond Paradox can be resolved if one allows for bounded rationality in the sense of noisy best-responding.

6.3 After the shock

Once a cost shock occurs subjects enter into a phase where the QRE is upset by the injection of asymmetric information. In order to explain the “Rocket-and-Feather” phenomenon of asymmetric price adjustment, we have to investigate how quickly a stable situation is reached in the different treatments. For this we look at the period when consumers have sufficiently updated their beliefs, such that a new QRE is reached. Ex post the estimated noise parameters can be used for some cross validation. If the noise parameters are similar to those we observed in the pre-shock phase then this is evidence for behavior having settled. Looking at how the price distributions evolved over time after the shock in the different treatments (Figure 4 and 5), suggests that in the *MC-Increase* treatment a once off shift happens right away after the shock, without significant changes in the price distribution later on. From our estimation of the search rule (Table 2) we see that the search rule also adjusts immediately after one period. This snap-adjustment seems to be driven by the updating done by buyers in the shock-period. Almost all buyers reacted after observing a higher price with searching (more than 80 percent of buyers searched), just to learn that prices are now generally high. If this story of adaptive expectations leading immediate adjustments is true we should observe a new QRE in the *MC-Increase* treatment already from period 17 on.

Figure 7 shows the empirical and QRE price distribution and the empirical and

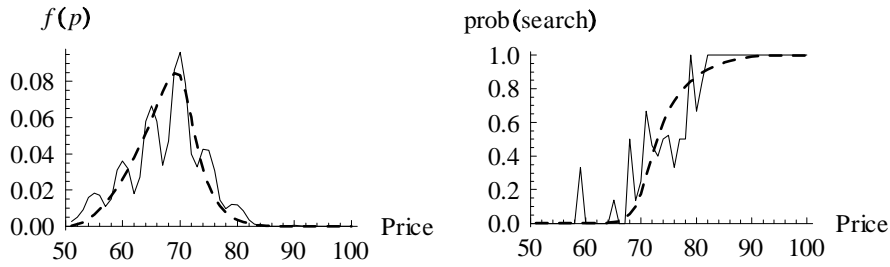


Figure 7: Post-shock QRE in the MC-Increase treatment

QRE search rule in the *MC-Increase* treatment for pooled data from one period after the shock to the final period. The noise parameters that achieved the best fit ($\mu_s = 0.5$ and $\mu_b = 0.21$) are remarkably close to those from the pre-shock phase.¹⁴ This shows that already from period 17 on (with only one updating period after the shock) a new QRE is reached.

As expected, we observe that in the treatment where the price-shock did not change the marginal cost (*MC-Constant*) the QRE did not shift. As the marginal cost had not changed, the firms kept their prices constant, adaptive consumers kept searching in the same way as before and not much changed until the end of the game. Figure 8 shows the QRE and the empirical price distribution and the search rule. The irrationality parameters are again close to the ones estimated for the pre-shock phase ($\mu_s = 0.67$ and $\mu_b = 0.27$) but slightly higher, whereas the fit is much improved.¹⁵ The slightly higher noise parameters stem from a somewhat more dispersed price distribution, which the participants in this treatment showed already in the pre-shock phase.

The picture is different for the *MC-Decrease* treatment. The price distribution is only stabilizing around period 24 (see Figures 1 and 2). Furthermore, the regression results for the search rule (see Table 2) show that there are large search spikes in periods 17, 20, and 23. So the search rule is also not stable in the periods before 24. So up to at least period 24 the adjustment to the cost shock is still ongoing and the resulting asymmetric information is not resolved yet. For this reason it is not

¹⁴The fit is even slightly better than in the pre-shock phase ($\chi^2 = 31.71$, same degrees of freedom; $d = 0.094$).

¹⁵We have $\chi^2 = 10.23$ (same degrees of freedom) and $d = 0.080$.




Figure 8: Post-shock QRE in the MC-Constant treatment

sensible to fit the QRE model to the whole after-shock period.¹⁶

Finally, we fit our model for the periods from 24 on to the end in the *MC-Decrease* treatment. This should give us some idea if the price and search-rule adjustment has taken place by then. Irrationality parameters in the vicinity of those in the pre-shock phase and the other treatments, would provide evidence for that. Figure 9 plots again the fitted and empirical price distributions and search rules. The noise parameters are again close to those in the other cases ($\mu_s = 0.43$, $\mu_s = 0.27$), which gives us some confidence in our claim that the adjustment has occurred by period 24. Also reassuring is the good fit, as one would expect that erratic behavior QRE cannot account for should take place less often in the final periods of the experiment.¹⁷




Figure 9: QRE in periods 24 to 30 in the MC-Decrease treatment

In this Section we have shown that the price dispersion arising from our exper-

¹⁶Just to confirm this, we fitted the QRE model for the whole period anyway. The fit of both the price distribution and the search rule is worse than in the other treatments and in the pre-shock phase. Furthermore, it seems only possible to either fit the price distribution or the search rule by varying the noise parameters. For a good fit of the search rule χ^2 gets very large (> 60), whereas minimizing χ^2 leads to a large d (> 0.11).

¹⁷We have $\chi^2 = 20.68$ and $d = 0.060$.

iments can be explained by QRE, a stochastic equilibrium concept assuming that players are boundedly rational and do not always choose best responses but play better responses with higher frequencies. Price-dispersion is a prerequisite for any price adjustment after a cost shock in our search environment. It also provides the necessary conditions for asymmetric price adjustment arising from adaptive updating by the consumers. We have also found evidence for our claim that the updating process after a positive cost shock happens immediately in the period after the shock, whereas due to the adaptive nature of the price-expectations of consumers updating takes longer when the shock is negative. Only after about seven to eight periods a new QRE occurs that has similar parameters as those in the other treatments and in the pre-shock phase. The observation that the pre-shock equilibrium virtually carries over to the phases after the shock, when the cost shock does not change the marginal cost, shows that firms are not trying to exploit their information advantage by increasing prices when the cost stay unchanged.

7 Conclusion

We implement the simplest search model we can think of (with two sellers and one buyer) in the laboratory to study asymmetric price adjustment to cost shocks. The environment is designed such that none of the typically important ingredients used to explain asymmetric price adjustment is present. In our environment, the standard theory predicts the occurrence of a unique monopoly-price equilibrium for all periods and treatments. Tacit collusion is ruled out, as the model already predicts monopoly prices, while price dispersion as an other potentially important driver of asymmetric cost shocks should not occur according to theory. Consequently, cost shocks should not have any influence and no price adjustment after a shock should occur at all. We observe persistent deviations from the equilibrium (i.e., price dispersion with prices well below the monopoly price) and asymmetric price adjustment (i.e., prices adjust immediately upwards after a positive cost shock, while the downward adjustment after a negative shock takes several periods). An analysis of individual behavior suggests that bounded rationality alone could drive the observed price dispersion,

which is the prerequisite for asymmetric price adjustment in related theoretical models. Quantal Response Equilibrium does well at explaining behavior, which suggests that bounded rationality is an important factor for asymmetric price adjustment. We also find evidence that, as suggested by theory, asymmetric price adjustment is driven by buyers' asymmetric learning process of the true cost after a cost shock. In our case the asymmetry in updating is not driven by heterogeneous buyers updating their rational expectations as suggested in most theoretical papers. Adaptive expectations of the buyers are driving the asymmetry: after a positive shock consumer search spikes (since prices have increased) and updating is immediate, while the lower prices after a negative shock reduce the search intensity. Sluggish updating allows the sellers to reduce the prices only gradually.

References

- Abrams, E., M. Sefton, and A. Yavas (2000). An experimental comparison of two search models. *Economic Theory* 16(3), 735–749.
- Bacon, R. W. (1991). Rockets and feathers: the asymmetric speed of adjustment of uk retail gasoline prices to cost changes. *Energy Economics* 13(3), 211–218.
- Ball, L. and N. G. Mankiw (1994). Asymmetric price adjustment and economic fluctuations. *Economic Journal* 104(423), 247–261.
- Baye, M. R. and J. Morgan (2004). Price dispersion in the lab and on the internet: theory and evidence. *RAND Journal of Economics* 35(3), 449–466.
- Borenstein, S., A. C. Cameron, and R. Gilbert (1997). Do gasoline prices respond asymmetrically to crude oil price changes? *Quarterly Journal of Economics* 112(1), 305–339.
- Braverman, A. (1980). Consumer search and alternative market equilibria. *Review of Economic Studies* 47(3), 487–502.
- Burdett, K. and K. L. Judd (1983). Equilibrium price dispersion. *Econometrica* 51(4), 955–969.

- Camerer, C. F. (2003). *Behavioral game theory: experiments in strategic interaction*. Princeton University Press.
- Carlson, J. A. and R. P. McAfee (1983). Discrete equilibrium price dispersion. *The Journal of Political Economy* 91(3), 480–493.
- Cason, T. N. and D. Friedman (2003). Buyer search and price dispersion: a laboratory study. *Journal of Economic Theory* 112(1549), 232–260.
- Chen, H. A., D. Levy, S. Ray, and M. Bergen (2008). Asymmetric price adjustment in the small. *Journal of Monetary Economics* 55(4), 728–737.
- David, D. D. and C. A. Holt (1996). Consumer search costs and market performance. *Economic Inquiry* 34(1), 133–151.
- Diamond, P. A. (1971). A model of price adjustment. *Journal of Economic Theory* 3(2), 156–168.
- Duffy-Deno, K. T. (1996). Retail price asymmetries in local gasoline markets. *Energy Economics* 18(1-2), 81–92.
- Dufwenberg, M. and U. Gneezy (2000). Price competition and market concentration: an experimental study. *International Journal of Industrial Organization* 18(1), 7–22.
- Eckert, A. (2002). Retail price cycles and response asymmetry. *Canadian Journal of Economics* 35(1), 52–77.
- Ellingsen, T., R. Friberg, and J. Hassler (2006). Menu costs and asymmetric price adjustment. Working Paper.
- Fischbacher, U. (2007). Z-tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics* 10(2), 171–178.
- Galeotti, M., A. Lanza, and M. Manera (2003). Rockets and feathers revisited: an international comparison on european gasoline markets. *Energy Economics* 25(2), 175–190.
- Goeree, J. K., C. A. Holt, and T. R. Palfrey (2002). Quantal response equilibrium and overbidding in private-value auctions. *Journal of Economic Theory* 104(1), 247–272.

- Goodwin, B. K. and D. C. Harper (2000). Price transmission, threshold behavior, and asymmetric adjustment in the u.s. pork sector. *Journal of Agricultural and Applied Economics* 32(3), 543–553.
- Grether, D. M., A. Schwartz, and L. L. Wilde (1988). Uncertainty and shopping behaviour: an experimental analysis. *The Review of Economic Studies* 55(2), 323–342.
- Hannan, T. H. and A. N. Berger (1991). The rigidity of prices: evidence from the banking industry. *American Economic Review* 81(4), 938–945.
- Karrenbrock, J. D. (1991). The behavior of retail gasoline prices: symmetric or not? *Federal Reserve Bank of St. Louis Review July/August(Jul)*, 19–29.
- Lewis, M. (2004). Asymmetric price adjustment and consumer search: an examination of the retail gasoline market. Working Paper.
- Luce, R. D. (1959). *Individual Choice Behavior: A Theoretical Analysis*. New York: Wesley.
- Mckelvey, R. and T. Palfrey (1998). Quantal response equilibria for extensive form games. *Experimental Economics* 1(1), 9–41.
- McKelvey, R. D. and T. R. Palfrey (1995). Quantal response equilibria for normal form games. *Games and Economic Behavior* 10(1), 6–38.
- Neumark, D. and S. A. Sharpe (1992). Market structure and the nature of price rigidity: evidence from the market for consumer deposits. *Quarterly Journal of Economics* 107(2), 657–680.
- Peltzman, S. (2000). Prices rise faster than they fall. *Journal of Political Economy* 108(3), 466–502.
- Reinganum, J. F. (1979). A simple model of equilibrium price dispersion. *Journal of Political Economy* 87(4), 851–58.
- Rob, R. (1985). Equilibrium price distributions. *The Review of Economic Studies* 52(3), 487–504.
- Stahl, Dale O., I. (1989). Oligopolistic pricing with sequential consumer search. *American Economic Review* 79(4), 700–712.

- Tappata, M. (2009). Rockets and feathers: Understanding asymmetric pricing. *RAND Journal of Economics* 40(4), 673–687.
- Tirole, J. (1988). *The Theory of Industrial Organization*. Cambridge: MA: MIT Press.
- Tsiddon, D. (1993). The (mis)behaviour of the aggregate price level. *Review of Economic Studies* 60(4), 889–902.
- Varian, H. R. (1980). A model of sales. *American Economic Review* 70(4), 651–659.
- Ward, R. W. (1982). Asymmetry in retail, wholesale, and shipping point pricing for fresh vegetables. *American Journal of Agricultural Economics* 64(2), 205–212.
- Yang, H. and L. Ye (2008). Search with learning: understanding asymmetric price adjustments. *RAND Journal of Economics* 39(2), 547–564.

A Experimental Instructions

(The instructions for all treatments are identical up to the show-up fee given (4, 6 or 8 Australian Dollars). With this variation we equalize earnings over different treatments, as cost changes alter the total surplus available.)

Thank you for participating in this experiment. Please read the instructions carefully. This is important, as understanding the instructions is crucial for earning money. Please note that you are not allowed to communicate with other participants during the experiment. If you do not obey to this rule we may exclude you from the experiment. If you have any questions, please raise your hand. We will come to answer your questions individually.

The currency in this game is called E-Dollars. At the end of the game we will convert the E-Dollars you have earned into real money. The exchange rate is **100 E-Dollars = 2 Australian Dollars**. You will also be paid a base payment of **AUD 6.00** for your participation.

- **Your task**

You will play a market game in this experiment. There are two types of players in the game: sellers and the buyers. You will be randomly assigned your role (either as a seller or a buyer) at the beginning of the experiment. Your role will be announced to you and fixed for the whole duration of the experiment. In each round we will randomly pair two sellers with one buyer. Each of the two sellers wants to sell one unit of a good which will cost the seller **MC = 30 E-Dollars** to produce and sell. The buyer wants to buy one unit of the good, which he values at **V = 100 E-Dollars**. The profits for seller will be the selling price minus the cost MC if a sale takes place and zero otherwise. The profit of the buyer will be the valuation V minus the selling price. Your task in this market is to make as high a profit as possible (, the higher your profit the higher is your monetary payout after the experiment).

- **The trading environment**

The game is composed of two decision-making stages: the sellers' stage and the buyer's stage. In the sellers' stage, the two sellers in the same group simultaneously set the prices in E-Dollars at which they want to sell. After both sellers have entered their selling prices, the buyers enter the game. In the buyer's stage, the buyer will be randomly given one out of the two prices offered by the two sellers in the group. Then the buyer can decide if he a) wants to buy from this seller at this price, or b) invest 15 E-Dollars to see the price of the other seller or c) to exit the market. In the case that the buyer decided to invest **C = 15 E-Dollars** to see the second price he can then decide a) to buy from the first seller, b) to buy from the second seller, or c) to exit.

- **Your Profit**

The round profits will depend on the prices set by the sellers and the buying and search decision of the buyers. Depending on the type (seller or buyer) the profits will be given as follows:

a) Sellers:

$Price(P) - cost(MC=30, initially)$	<i>if the unit was traded</i>
<i>zero</i>	<i>if the unit was not traded</i>

Note that the production cost MC is only incurred if the unit is actually traded. Furthermore, the production cost is initially fixed at 30 but may change during the game (see below).

b) Buyers:

Valuation($V=100$)-Price(P)-searchcost if the unit was purchased

zero-searchcost if the unit was not purchased

Note that the search cost is $C = 15$ if the buyer invested in seeing the second price and zero if he did not.

• Repetition and cost shocks

You will play 30 rounds of this game in succession. You will always play the same role (buyer or seller); but you will play with changing partners in your group. The groups are newly and randomly formed after each round.

Recall, that the seller initially has production cost MC of 30. This cost stays the same for the first 15 rounds. In between rounds 15 and 16 there might be a cost shock (MC might take on a different value). Then the remaining rounds (16 to 30) will be played with the new cost. You will be given details about the cost shock on the screen once it occurs.

• Summary

In this market game you will be a buyer or a seller. If you are a seller you want to sell a unit of a good, if you are a buyer you want to buy a unit of the good. There are always two sellers setting prices simultaneously. They are paired with one buyer, who does only observe the price of one of the sellers initially. The seller of which a buyer sees the price is randomly determined. Then the buyer can decide to buy from this seller or to spend some search cost in order to learn the price of the second seller before making a purchasing choice.

Production costs are initially fixed at $MC=30$ and might change between rounds 15 and 16. This will be announced between rounds 15 and 16.

Again, please make sure that you understand the instructions clearly, as this is crucial for your earnings in this experiment. If you have any questions please raise

your hand. We will come and answer your question. Once you are ready, we will play a trial period, which is of no consequence for your payout, after which you can raise your hand again and ask questions before we start with the 30 rounds, which will determine your earnings.