universität wien

DIPLOMARBEIT

Titel der Diplomarbeit

Asymmetric Pricing due to Tacit Collusion in a Simple Dynamic Oligopoly Model

Verfasser

Martin Obradovits

angestrebter akademischer Grad

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

Wien, im April 2010

Studienkennzahl lt. Studienblatt:	A 140
Studienrichtung lt. Studienblatt:	Diplomstudium Volkswirtschaft
Betreuer:	Dr. Paul Pichler

Acknowledgments

First and foremost, I want to thank my parents Idella and Josef Obradovits for their extensive support. Without their love, encouragement and open-mindedness, I wouldn't be where I am now. In addition, my parents contributed to the quality of this thesis during many interesting discussions. Finally, their generous financial support enabled me to fully concentrate on my studies throughout the years.

I am sincerely thankful for my girlfriend Kathrin Gruber, who gave me a shoulder to lean on and emboldened me whenever I had problems. Just as important, Kathrin made countless helpful suggestions concerning the layout and structure of this work. Although being busy herself, she always had an open ear for my thoughts associated with writing this thesis.

Last but not least, I want to address special thanks to my excellent supervising tutor Paul Pichler. This thesis wouldn't have been possible without his experience and the exhaustive help and criticism he provided. I am especially grateful for all of his suggestions concerning the model of asymmetric pricing I invented. Furthermore, Paul's quick response to questions and his continuous feedback was very beneficial.

Contents

1 Introduction 1

2 Types of Asymmetric Pricing 6

3 Asymmetric Pricing due to Tacit Collusion in a Simple Dynamic Oligopoly Model 11

3.1 Intuition and Model Setup 12

3.2 Equilibrium of the Stage Game 16

3.3 Equilibrium of the Dynamic Game 19

3.4 Asymmetric Pricing with Separated Submarkets 35

3.5 Summary and Discussion 42

4 Asymmetric Price Adjustment due to Consumer Search 45

4.1 A Model of Consumer Search With Learning 46

4.2 A Model of Consumer Search Without Learning 64

4.3 Other Consumer Search Models 78

4.3.1 A Reference Price Search Model with Asymmetric Pricing 78

4.3.2 Asymmetric Pricing caused by Sticky Consumer Prices 82

5 Miscellaneous Sources for Asymmetric Price Adjustment 85

5.1 Asymmetric Pricing due to Lags in Adjustment of Production and Finite Inventories 86

5.2 Asymmetric Pricing due to Asymmetric Menu Costs 87

5.3 Edgeworth Price Cycles Resembling Asymmetric Pricing . . 88

6 Summary and Conclusion 93

Bibliography **95**

List of Figures **97**

Appendix A **99**

A.1 Abstract 99

A.2 Abstract (German) 100

A.3 Curriculum Vitae 101

1 Introduction

This thesis studies the phenomenon of *asymmetric pricing*, according to which output prices tend to adapt differently to input cost increases than to input cost decreases. Empirically, asymmetric pricing (aka *asymmetric price adjustment* or *rockets and feathers*[1]) is typically associated with negative cost shocks being passed along to consumers in a 'slower' fashion than positive cost shocks. The general public as well as government authorities often attribute this pricing behavior to an abuse of market power, i.e., implicit or explicit collusion. However, to the best of my knowledge, no formal model capable of generating rockets and feathers through collusion has been developed. Therefore, the main contribution of my thesis will be to provide a collusive model of asymmetric price adjustment. Furthermore, I will summarize several other models of asymmetric pricing that do not rely on an abuse of market power and compare their results.

Why is asymmetric pricing even a relevant topic, and which support exists for it? During the last two decades, numerous empirical studies confirmed the so called pattern of rockets and feathers. Because of the rich amount of data available and some theoretical reasons that make gasoline a good candidate to test for asymmetric pricing, many studies have concentrated on the gasoline market. Examples include Johnson (2002); Bacon (1991); Karrenbrock (1991); Borenstein et al. (1997) and Lewis (2009). However, there is growing support that asymmetric pricing is not restricted to few specialized markets, but is a very broad phenomenon. Yang and Ye (2008, p. 547) state that other markets where adjustment asymmetry can be found include

[1] Output prices rise like rockets when there is a positive cost shock, but fall like feathers when there is a negative one.

> fruit and vegetables (Pick et al., 1991; Ward, 1982), beef and pork (Boyd and Brorsen, 1988; Goodwin and Holt, 1999; Goodwin and Harper, 2000) and banking (Hannan and Berger, 1991; Neumark and Sharpe, 1992; O'Brien, 2000).

Finally, in his comprehensive study of 77 consumer and 165 producer goods, Peltzman (2000) shows that asymmetric pricing can be found in more than two thirds of the markets he observed. He also points out that positive cost shocks usually have twice the immediate impact on output prices than negative ones, and that this asymmetry tends to last for at least five to eight months.

But while asymmetric price adjustment is found to be common in a great number of markets, it is not well understood from a theoretical perspective. As a consequence, modern economic theory seems incomplete and the field of asymmetric pricing could turn out to greatly enhance our understanding of how markets work. As Peltzman (2000, p. 468) points out:

> If [asymmetric pricing] was shown to be general and not just limited to a few case studies, it would point to a serious gap in a fundamental area of economic theory.

So which ideas and models are available that can lead to a better understanding of the phenomenon? First of all, Borenstein et al. (1997) provide three hypotheses which could explain adjustment asymmetry in gasoline retail. Some of the ideas mentioned in their paper might easily translate to other markets. Probably most interesting is the first hypothesis, which states that

> [p]rices are sticky downward because when input prices fall the old output price offers a natural focal point for oligopolistic sellers. (Borenstein et al., 1997, p. 324)

In brief, firms in oligopolistic markets could choose to maintain the price charged before a negative cost shock, as this price provides a focal point for (implicit) collusion. The idea for this hypothesis stems from the so called 'trigger sales' model of Tirole (1988), however, the economic literature still seems to lack an explicit formal analysis of this mechanism.[2]

The principal contribution of my thesis will thus be to try to overcome this theoretical deficit by providing a simple model of asymmetric pricing caused by tacit collusion in a trigger sales context. This will be the focus of the main section of my thesis. There, I will explain the basic mechanism that leads to asymmetric price adjustment in my model, describe its setup, prove that a collusive equilibrium exists for certain parameter values and present several comparative statics results. I will then extend the basic model to a multitude of separated submarkets and show that this model can reproduce realistic pricing patterns. Finally, I will point out some possible extensions that could be considered for future research.

In the rest of my thesis, I will explain most other models of asymmetric price adjustment that are currently available. Which models are these? First of all, the other proposed hypotheses for asymmetric price adjustment in Borenstein et al. (1997) were production lags with finite inventories and a signal extraction problem for consumers when prices are very volatile in a consumer search scenario. While

[2] Support for this can be found in Lewis (2009, p. 18), who writes that

> No rigorous model of focal price collusion has been specified, and testing against the predictions of a super-game model of tacit collusion is difficult since there are an infinite number of equilibrium price paths.

and Eckert (2002, p. 53), who states that

> While a theory linking tacit collusion with an asymmetric response to input prices has not been formally derived, the intuition has been discussed in the context of the Green and Porter (1984) model of tacit collusion with price wars.

the first explication has not been examined in detail, consumer search has become *the* number one modeling approach for rockets and feathers.
In fact, at least four rigorous models capable of generating 'true' rockets and feathers through consumer search exist. They have been invented by Yang and Ye (2008); Tappata (2009); Lewis (2009) and Cabral and Fishman (2008). As one can see, these papers are pretty recent, with the latter two being unpublished as of now. Although consumer search is the main economic concept which drives them, they build on different assumptions and provide heterogeneous mechanisms to derive adjustment asymmetry of output prices. I will provide a detailed summary and comment on limitations and possible extensions of the first two models in Section 4 of my thesis. There, I will also give a brief intuition of how the latter two models work.

Alternatives to consumer search include the mechanism of production lags with finite inventories proposed by Borenstein et al. (1997) as well as the models of Ball and Mankiw (1994) and Eckert (2002). Ball and Mankiw create a model of asymmetric pricing through asymmetric menu costs in inflationary environments. However, their model comes with several conceptual problems and is unlikely to be relevant in most markets. Eckert provides a model of Edgeworth price cycles that results in a pattern of price changes resembling asymmetric pricing although prices move independently from costs. Interestingly, an extension of this model could lead to true asymmetric pricing. See Section 5 for a further discussion of these three approaches.

The thesis will be organized as follows. In the next section, I will try to give a more formal definition of asymmetric pricing, both listing all the relevant types of rockets and feathers and providing graphical examples. In the main section (Section 3) of this work, I will derive and discuss an own model of asymmetric price adjustment in a collusive setting. In Section 4 and Section 5, I will proceed to present several other ideas how to model the phenomenon. I will also discuss problematic aspects and possible extensions of these ideas. A short summary (Section 6) concludes.

2 Types of Asymmetric Pricing

Before I will start to present various models of asymmetric pricing, it seems useful to give the reader a better understanding of how the phenomenon can be understood. In order to do so, I will distinguish between the various types of rockets and feathers discussed in the literature and provide a formal definition. Several diagrams are included for convenience.

The definition of asymmetric pricing given is based on Karrenbrock's (1991) definition of asymmetric gasoline price movements and the general definitions of Meyer and Cramon-Taubadel (2004, p. 538 ff.). According to the former,

> [r]etail price movements are defined as asymmetric if an increase in the wholesale price affects the retail price differently than an equal-sized decrease. (Karrenbrock, 1991, p. 22)

Karrenbrock then proceeds to list three distinct types of asymmetric price adjustment. The first type deals with the speed of adjustment. If an input price increase is passed along more quickly than an input price decrease, he refers to *time asymmetry*. A typical example for time asymmetry would be an immediate and full pass-through of positive cost shocks, but a delayed, full pass-through of negative cost shocks. In other words, an equal amount of some equal-sized, positive or negative input price change is passed along to output prices eventually, but adjustment needs longer when prices have decreased. The following figure taken from Meyer and Cramon-Taubadel (2004) illustrates this first type of adjustment asymmetry at the top panel, with p^{in} referring to some upstream input price and p^{out} referring to the output price. The shaded area represents the welfare loss that is implied for consumers.

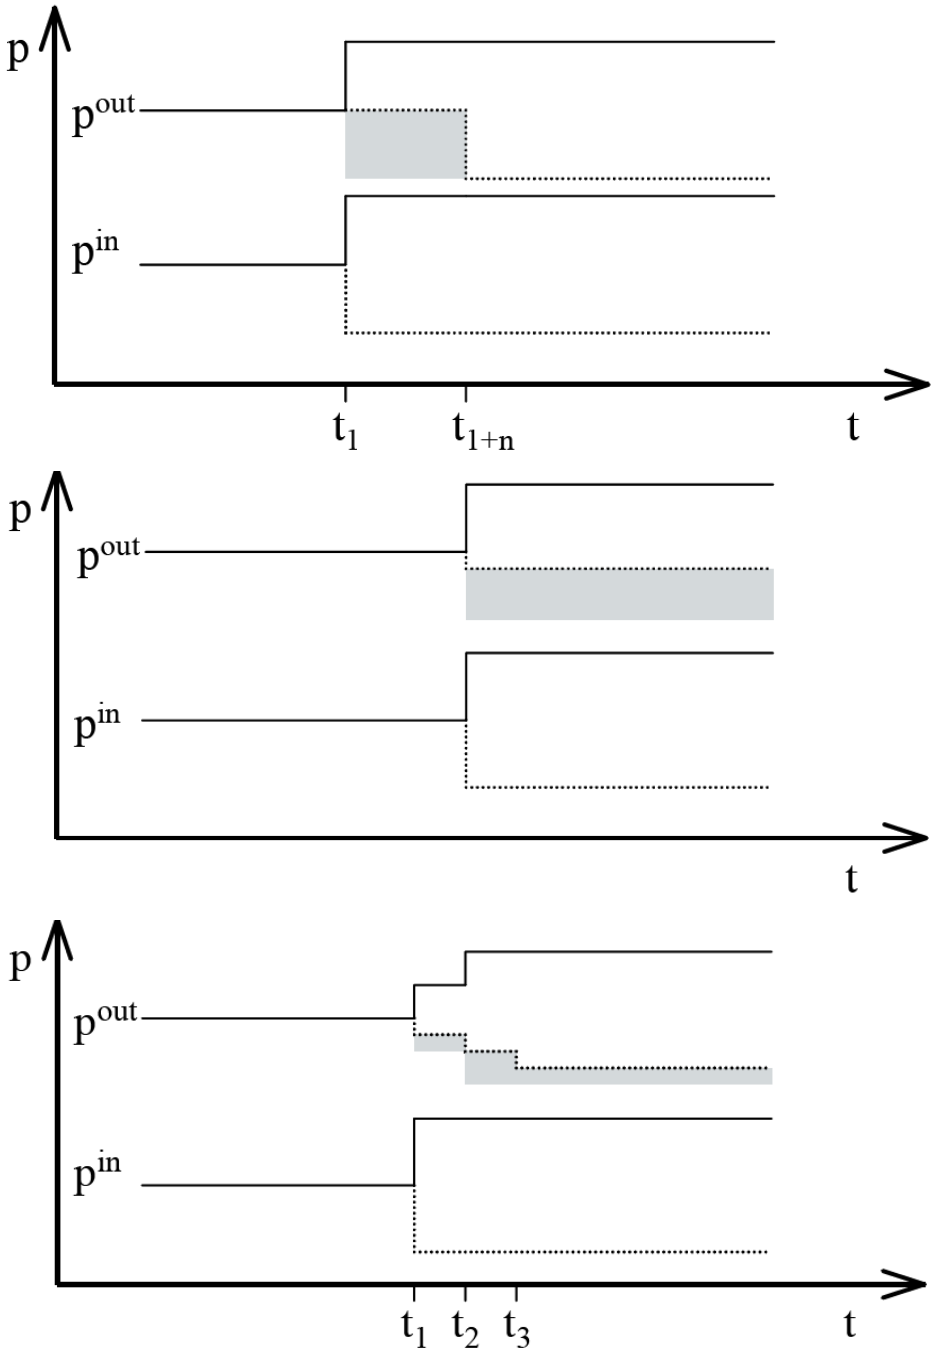

Figure 2.1 Time, amount and pattern asymmetry. Source: Meyer and Cramon-Taubadel (2004, p. 584)

The second type of asymmetry is called *amount asymmetry.* This means that, once the adjustment process is over, a higher percentage of an input price increase will be reflected in the new output price, compared to an input price decrease. In this case, price increases and decreases induce an equally long adjustment process, but after the phase of transition, output prices will have changed more (in absolute terms) after a positive cost shock than after a negative one. For example, if after two weeks an input price increase of ten cents results in an output price increase of eight cents, while an input price decrease of ten cents only results in an output price decrease of six cents, this can be referred to as amount asymmetry. A graphical example can be found in the middle panel of Figure 2.1.
It is worth noting that true amount asymmetry is unlikely to prevail in actual markets. The reason is that input and output prices are usually not independent, but affect each other through some kind of equilibrium relation. But true amount asymmetry implies that prices would drift apart over time, destroying any projected equilibrium relation. In other words, amount asymmetry can only occur over the short run, with the long run equilibrium being restored over time. (Wikipedia, 2009)

The final and most important type of asymmetry is a combination of the types explained above. Price adjustment can differ in its length *and* amount. For example, it might be the case that an input price increase leads to a full output price increase after one day, while an input price decrease needs three days to be transmitted, with only 80% of the price change being finally reflected in the new output price.
But even if both the length and amount of price adjustment are equal, the *patterns* of the adjustment process can differ. For example, a three-week full price adjustment after an input price increase of ten cents could be (5, 3, 2) cents,

while it might be (2, 4, 4) cents when facing an input price decrease. As this is also clearly a case of adjustment asymmetry, the third type will be referred to as *pattern asymmetry* and includes the combination of time and amount asymmetry mentioned first. For a graphical example, see the bottom panel of Figure 2.1.

The above illustrations might induce the thought that amount asymmetry (or amount asymmetry combined with time asymmetry, but not pure pattern asymmetry) is the worst type of adjustment asymmetry, as consumers would have to bear 'infinite' welfare losses. This is obviously not the case, as the net present value of the welfare loss induced by large time asymmetry might very well exceed the net present value of the loss induced by small amount asymmetry that persists forever.

As a final note, it is important to notice that all of my examples refer to a state called *positive asymmetric price transmission* (PAPT), which simply means that output prices react more quickly (or completely) to input price increases than to input price decreases. PAPT is the most commonly observed type of adjustment asymmetry and is the sole focus of my thesis. For the sake of completeness, the opposite case is called *negative asymmetric price transmission* and implies that output prices tend to adjust more rapidly to input price decreases than to increases, i.e., that there is a welfare redistribution from producers to consumers. Figure 2.2 showcases this behavior.
A last distinction could also be made between vertical and *spatial* asymmetric price transmission, but in my thesis, only vertical relations will be covered.

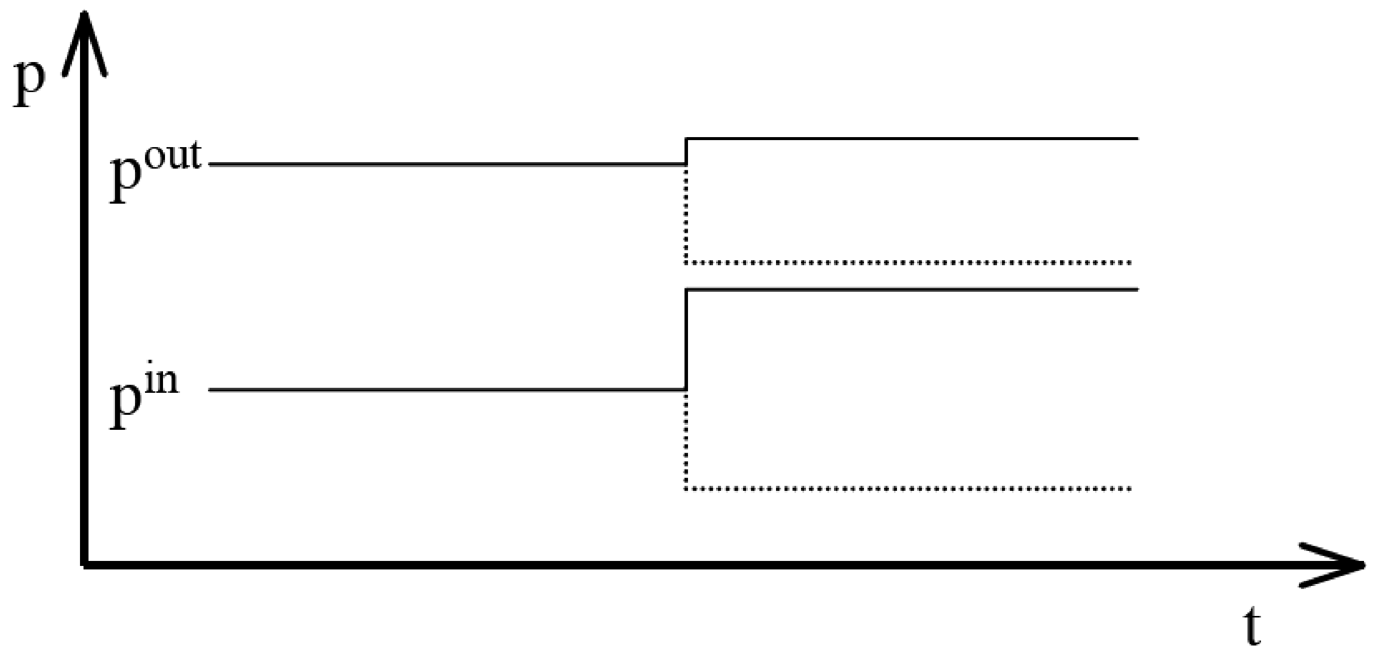

Figure 2.2 Amount asymmetry with negative asymmetric price transmission. Source: Meyer and Cramon-Taubadel (2004, p. 586)

3 Asymmetric Pricing due to Tacit Collusion in a Simple Dynamic Oligopoly Model

In this main section of my thesis, I develop a model of asymmetric pricing based on tacit collusion among firms. The idea to this stems from the first hypothesis of how asymmetric might emerge proposed in Borenstein et al. (1997). There, the authors argue that asymmetric pricing might occur because firms who compete in Bertrand competition coordinate their selling prices on the price that was charged before a negative cost shock, but have to immediately increase their prices after a positive cost shock because margins are squeezed and firms would have to incur losses otherwise. Due to random demand shocks, collusion slowly breaks down after a cost decrease, leading prices to slowly adapt to negative cost shocks. In contrast, positive cost shocks are transmitted without delay to output prices.

This section will be organized as follows: In Subsection 3.1, I will give a more detailed intuition of how asymmetric pricing emerges in the model and will outline its setup. In Subsection 3.2, I will calculate the equilibrium of the static game, i.e., the equilibrium of the game when it is not repeated. In Subsection 3.3, I propose a collusive strategy combination that turns out to be an equilibrium of the dynamic game. Doing so, I will provide several comparative statics and ultimately show that asymmetric pricing results for certain parameter values. Next, in Subsection 3.4, I will show that a more realistic pattern of price adjustment can be found if one thinks of one market to be composed of several separated submarkets, as has already been proposed in Borenstein et al. (1997). A short summary and discussion (Subsection 3.5) concludes.

3.1 Intuition and Model Setup

The basic mechanism which leads to asymmetric pricing in my model is as follows. After a negative cost shock, firms witness rising margins. While there could be fierce competition and a quick negative response of the output price to the cost shock, firms can also collude on the price charged before the cost shock because it provides a natural focal point for collusion.

In traditional models of industrial organization, firms can observe the prices of all their competitors[3] and thus, depending on some discount factor δ, only one type of collusive equilibrium is possible. The strategies involved in this equilibrium are to price at the collusive price whenever all of the other firms price at the high price, and to price at cost forever (or at least several periods) as punishment if any firm deviates. In this scenario, either all firms collude forever on some high price that maximizes their profits or all firms price at cost (deviation from the collusive behavior either pays or not, so firms either have no incentive to deviate in every period or no collusive equilibrium exists).

However, things change drastically if one introduces imperfect information. If firms cannot observe the prices of their rivals, they would have to find other mechanisms to keep their competitors from deviating. One way to do so is by making collusion dependent on sales in previous periods. I will call a strategy that conditions cooperation on previous sales a *trigger sales strategy*. The intuition behind such a strategy is that firms can infer from low sales that other firms have broken the tacit collusive agreement and will end collusion as well. From this, it is only a

[3] Usually facing a delay of one period.

small step to asymmetric pricing. If firms confuse random demand shocks with changes in demand caused by firms that deviate from the collusive price[4], collusion can slowly break down due to random demand shocks although every firm wishes to collude forever.

In contrast, positive cost shocks are translated without delay because margins are very small and would become negative otherwise. The result is asymmetric price transmission: positive cost shocks are transmitted immediately to output prices whereas negative cost shocks need some time to have a negative effect on prices.

Since demand is random, the source of randomness is crucial for determining the speed with which collusion breaks down, and how prices react once collusion ends. In the easiest case, only overall demand is random and all firms end collusion at the same time in symmetric equilibria, i.e., when sales fall below some critical value that is defined by a trigger sales strategy.[5] Because of this, all prices in the market would have to drop sharply and simultaneously once collusion ends. This behavior is usually not observed in real markets. However, a possible solution to

[4]For this to be possible, there must be some fraction of uninformed consumers that buy at a random firm because otherwise, firms can always accurately determine whether some competitor has deviated to a lower price.

[5]Another possibility would be that demand is random for every firm and also overall demand is random. Unfortunately, it seems that no trigger sales equilibrium can be found for this model. The reason for this is that if every firm followed a strategy of deviating when demand falls under some critical value, a best response to this strategy combination would be to always price at the collusive price, implying higher profits (if all other firms do not deviate, a firm shouldn't deviate if it faces a random demand shock). Although there are certainly ways to overcome this issue by introducing beliefs or other sophistications, I will leave this open for future research.
Also, a two-firm model can be conceived where overall demand is constant, but each firm faces random demand (i.e., overall demand is split randomly among both firms). In this case, the strategy combination of playing the collusive price as long as sales do not fall below some threshold *and* sales of the other firm do not fall below the same threshold (which, in this case, can be inferred from the firm's own demand) results in an equilibrium. While this would again imply asymmetric pricing, I will not discuss a model of this type in my thesis.

this problem is if one thinks of one big 'market' to be comprised of several independent submarkets. This assumption is probably justified for many 'markets' that are in fact composed of spatially separated submarkets (e.g. in gasoline retail, a city could be defined as one market although there is an inner city market, suburban market, highway market, etc. with all of those being independent from each other). In this case, collusion can progressively break down in parts of the 'market' (each with random demand), which leads to a smooth decline of average market prices.

First, I will provide a model of the former type, i.e., a model where collusion is maintained for several periods and then suddenly collapses, leading to a full and abrupt decline of average market prices. Then, I will extend this model to allow for separated submarkets, resulting in a more realistic pattern of price response to negative cost shocks.

For the former, consider a market with $n \geq 2$ firms who produce a homogeneous good at marginal cost c. Firms compete à la Bertrand and try to maximize expected profits over the current and future periods. Moreover, while firms can perfectly observe the demand they faced in every elapsed period, they can never observe the prices charged by their competitors (both for current and bygone periods).

Time is discrete, with $t = 1, 2, 3, ...$ The interval between two periods is assumed to be short (e.g. one day) such that firms' discount factor is close to one. For simplicity, I will suppose that firms do not discount over future periods.[6] Each

[6]In contrast to classic trigger strategy games, the model doesn't rely on discounting because its structure ensures that different cost states will have a finite expected length. The model is robust to the introduction of a discount factor as long as it is sufficiently close to one. For a short interval between periods, this is a realistic assumption.

period, the n firms in the market have to pay either low or high marginal cost which is common to all firms, $c \in \{c_L, c_H\}$, with $c_L < c_H$. Cost states follow a two-state Markov process with transition probability $(\rho, 1-\rho)$, with $\rho > 1/2$. Put differently, it holds that $\Pr(c^t = c_L | c^{t-1} = c_L) = \Pr(c^t = c_H | c^{t-1} = c_H) = \rho > 1/2$.

The consumer side is characterized by a continuum of (potential) consumers with random measure $X \sim U[0,2]$[7] that is identically and independently distributed (i.i.d.) in each period, with $E[X] = 1$. This random measure is unobservable for firms. Also, consumers have a reservation price ν that is equal to c_H, i.e., $\nu = c_H$, and unit demand in each period. A fraction $\lambda \in (0,1)$ of consumers is informed and observes all prices in the market. These consumers are called *shoppers* and buy at the lowest priced firm (if two or more firms have the same lowest price, shoppers buy at one of the firms at random). The remaining fraction $1 - \lambda$ of nonshoppers only observes one random price and thus buys at a random firm. Consumers do not buy if a price exceeds their reservation price ν. Note that my model is a partial equilibrium model: while firms behave optimally, consumers can only decide whether to buy or not. For models where the search decision of consumers is endogenous, see Section 4.

The parameters n, λ, ρ, c_L, c_H, ν as well as the distribution of potential consumers are assumed to be common knowledge.

Because of the above assumptions, the expected demand of some firm i, pricing at $p \leq \nu$, can be written as

[7]The uniform distribution is mainly chosen for technical convenience because it has a simple and closed form cumulative distribution function. In principle, every continuous distribution with a bounded or semi-infinite interval that has zero as lower bound can be chosen. If minimum market demand is greater than zero, it depends on the model parameters whether a collusive equilibrium can exist.

$$E(D_i|p) = \begin{cases} E[(\frac{\lambda}{b} + \frac{(1-\lambda)}{n})X] = \frac{\lambda}{b} + \frac{1-\lambda}{n} & |\text{if } p = p_{min} \\ E[(\frac{(1-\lambda)}{n})X] = \frac{1-\lambda}{n} & |\text{if } p > p_{min} \end{cases} \quad (3.1)$$

where $p_{min} = min\{p_1, ..., p_n\}$ and a total of $1 \leq b \leq n$ firms price at p_{min}.

In what follows, I will derive Nash-equilibria for the stage game[8] and, building on this, the dynamic game where the stage game is repeated infinitely. Once equilibrium for the dynamic game in the market has been derived, I will proceed to analyze the effect on average market prices when there are multiple separated submarkets.

3.2 Equilibrium of the Stage Game

If the above game is only played for one period (i.e., only for $t = 1$), the strategy space of firms is restricted to either price at some level for sure (playing a pure strategy) or to randomize between prices (playing a mixed strategy).
It is straightforward to see that there can be no symmetric equilibrium in pure strategies if $c = c_L < \nu$. The reason for this is an undercutting argument. Suppose every firm played some pure strategy $p \in (c_L, \nu]$. Then each firm would have an incentive to slightly undercut all of its rivals because it can increase its expected demand without effectively decreasing its profit margin. This goes on until every firm would price at c_L, driving profits to zero. But then again, each firm would have an incentive to price at some higher price and make a positive profit by getting positive expected demand of $\frac{1-\lambda}{n}$ from the nonshoppers. Therefore, no equilibrium in pure strategies can exist if costs are low.
On the other hand, if $c = c_H = \nu$, in principle every price that is greater or equal to the valuation of consumers ν can be played as equilibrium strategy, each resulting

[8]For this, I will use a methodology similar to Varian (1980).

in a profit of zero. I will restrict the action space of firms such that they cannot price above ν.[9] Doing so, the only equilibrium of the stage game when costs are high is that every firm prices at ν. This is a crucial feature of the model: only because firms do not randomize over prices under high costs, there exists a unique focal point for collusion (i.e., $p = \nu$) once costs drop.

Because of the above result that there can be no pure strategy equilibrium of the stage game if costs are low, in any symmetric low cost equilibrium, firms will have to price using a probability distribution $F(p) := \Pr(\tilde{p} \leq p)$ with support $[\underline{p}, \overline{p} = \nu]$.[10] The trade-off between attracting shoppers and extracting high profits from non-shoppers is resolved by using mixed strategies. In equilibrium, *surplus-appropriation* must be balanced by *business-stealing* effects, which is only possible when firms randomize between prices (see Varian, 1980; Tappata, 2009, p. 677). The equilibrium price distribution can now be calculated by the following logic: playing the mixed strategy of $F(p)$ can only be optimal if any price in the support $[\underline{p}, \overline{p} = \nu]$ yields the same expected profit. Otherwise, the mixed strategy $F(p)$ could profitably be altered by choosing prices that generate lower profits less frequently. Since ν is clearly in the pricing support of firms, it must hold in particular that setting any price p in the support has to yield the same expected profit as setting $p = \nu$. Formally, for $F(p)$ to be an equilibrium strategy it must hold that

$$E(\Pi_i(p; F(p)) = E(\Pi_i(v; F(p)), \tag{3.2}$$

[9] This can be motivated if one assumes firms to strictly prefer positive demand over zero demand.

[10] ν has to be the highest price in the support because it would make no sense for a firm to price at $\overline{p}$ if $\overline{p} < \nu$: the firm would not attract any informed consumers anyway and could thus make a higher profit by pricing at ν.

where $\Pi_i(x; F(p))$ denotes the profit of firm i if it prices at x, given that all other firms in the market price according to the probability distribution function $F(p)$.

Using the expected demand function of firms given by Equation (3.1) and realizing that $b = 1$ when all other firms price according to $F(p)$ (which has no mass points[11]), the left hand side of Equation (3.2) is given by

$$\begin{aligned}
E(\Pi_i(p; F(p)) &= (p - c_L)[E(Di|p = p_{min})\Pr(p = p_{min}) + E(Di|p > p_{min})\Pr(p > p_{min})] \\
&= (p - c_L)\left[\frac{\lambda n + 1 - \lambda}{n}(1 - F(p))^{n-1} + \frac{1-\lambda}{n}(1 - (1 - F(p))^{n-1})\right] \\
&= (p - c_L)\left[\lambda(1 - F(p))^{n-1} + \frac{1-\lambda}{n}\right].
\end{aligned}$$

From this it immediately follows that the right hand side of Equation (3.2) can be written as

$$E(\Pi_i(\nu; F(p)) = (\nu - c_L)\frac{1-\lambda}{n}.$$

Setting both sides equal and rearranging finally yields

$$F(p) = 1 - \sqrt[n-1]{\left(\frac{\nu - p}{p - c_L}\right)\frac{1-\lambda}{\lambda n}}. \tag{3.3}$$

Using that $F(\underline{p}) = 0$, one can also solve Equation (3.3) for $\underline{p}$. It holds that

$$\underline{p} = \frac{\nu(1-\lambda) + \lambda n c_L}{1 - \lambda + \lambda n} > c_L. \tag{3.4}$$

[11] A rigorous proof for this can be found in Varian (1980).

One can see that $\underline{p}$ will only reach c_L if $\lambda = 1$ (i.e., every consumer is informed) or $n \to \infty$. Interestingly, the equilibrium of the stage game collapses to a full competition Bertrand equilibrium when $\lambda = 1$, but ends up in the monopoly case (with each firm pricing at ν) when there is an 'infinite' number of firms in the market. This is because the expected profit from business-stealing decreases with a higher rate than the expected profit from surplus-appropriation for increasing n, i.e., having the lowest price in the market will pay less and less for higher n. Of course, the game will also collapse to the monopoly outcome if the fraction of informed consumers λ is zero.

3.3 Equilibrium of the Dynamic Game

Now that the equilibrium strategies for the static game have been determined, I will proceed to propose a collusive equilibrium for the dynamic game, i.e., the stage game when it is repeated infinitely.

First of all, it is straightforward to see that the strategy combination of pricing at $p = \nu$ whenever $c = c_H$ and pricing according to $F(p)$ whenever $c = c_L$ clearly constitutes a subgame perfect Nash-equilibrium of the supergame. This is because by definition, the equilibrium of the stage game must be an equilibrium of *every* stage (period). Thus, if *other* equilibria for the dynamic game exist, it is in principle impossible to pin down which equilibrium will be played. In fact, the well known *Folk theorem* states that in repeated games, *any* strategy combination can constitute an equilibrium as long as each player's *minimax* condition is fulfilled, meaning that each player minimizes their maximal loss. (See Wikipedia, 2010, and the references therein)

More specifically, if costs drop and firms want to collude on some high price to keep their margins at a high level, there are generally infinitely many prices firms can collude on in equilibrium, rendering it unclear how they can coordinate on *one* particular price. Because of this, I will only examine equilibria where firms' collude on the price that was charged before a negative cost shock, as it provides a natural focal point for collusion. It seems plausible that, if there is no other information available, firms' best guess would be to keep prices unchanged in order to coordinate their selling prices after a cost decrease.

Also, among all possible collusive equilibria, I will only consider *trigger sales* equilibria in which firms' profits are maximized. As will turn out below, this implies that, depending on the parameters n and λ in the market, either the equilibrium of the stage game will be played in every period (I will call such markets 'fierce', as firms engage in fierce competition) or a collusive trigger sales equilibrium emerges (I will call such markets 'collusive').

I will now prove that, for some values of n and λ, the following trigger sales strategy combination constitutes a subgame perfect Nash-equilibrium:

- Price at $p = \nu$ whenever $c = c_H$.
- Price at $p = \nu$ in the first period after costs drop from c_H to c_L.
- Price at $p = \nu$ in subsequent periods where $c = c_L$ as long as the demand faced in every elapsed period (where $c = c_L$) exceeded a critical value k, with $0 < k < 2\left(\frac{1-\lambda}{n}\right)$.[12] If demand has been lower than k in some period, price according to the equilibrium of the low cost static game (i.e., price according to $F(p)$) until costs rise to c_H again.

[12] A minimum demand of zero (implying that firms always collude) can never be an equilibrium, as will be outlined below. $2\left(\frac{1-\lambda}{n}\right)$ is the maximum demand (demand if $X = 2$) a firm can face if some other firm has deviated by undercutting. Because of this, a strategy of breaking collusion conditional on having witnessed a demand greater or equal to this value (i.e., $k \geq 2\left(\frac{1-\lambda}{n}\right)$) makes no sense. This is because for $X \in [2\left(\frac{1-\lambda}{n}\right), k]$, firms would break collusion as punishment even though they would know with certainty that every other firm in the market colluded in the previous period.

As will be shown, there is generally a continuum of values of k (and therefore infinitely many trigger sales strategies) that result in an equilibrium of the dynamic game. However, as mentioned above, I will determine k endogenously by assuming that firms collude on the value of k that maximizes their expected profits while still keeping firms from deviating.

Which value of k is that? In a world with perfect information, firms would observe overall market demand and could therefore accurately deduct whether low demand is caused by some rival firm undercutting or a random demand shock. However, in the model firms do not get to know the random demand (which ranges from 0 to 2) and will thus have to break collusion if demand falls below some critical value $k > 0$ even if every firm wishes to collude. Otherwise, it would pay to deviate because if there is no punishment for deviation, it must always be profitable to do so: all other firms would keep playing ν anyway no matter how low their demand gets.

The optimal (profit maximizing) value of k is thus given by the minimal k (say k^*) that is still incentive compatible with not deviating from the collusive strategy. This is because the lower k, the less likely it is that market demand falls below this threshold value despite of every firm colluding.

Having outlined the intuition behind the critical value k, some preliminary results are needed to determine combinations of the number of firms n, persistence parameter ρ, number of informed consumers λ and critical value k that allow for a collusive equilibrium.

I will denote by q the probability that demand reaches at least k for any firm playing ν if all other firms play ν. As the measure of potential consumers is random, $X \sim U[0,2]$, the demand for such a firm is random as well, with $D_i = (\frac{\lambda}{n} + \frac{(1-\lambda)}{n})X = \frac{X}{n}$. Now $q := \Pr(D_i \geq k|\nu;\nu) = \Pr(\frac{X}{n} \geq k) = 1 - \Pr(X < nk)$. Using the distribution function of a uniformly distributed random variable with support $[0,2]$, it follows that

$$q = 1 - \frac{nk}{2} \tag{3.5}$$

if $0 < k < 2\left(\frac{1-\lambda}{n}\right)$, as assumed.

Analogously, I will denote by r the probability that demand reaches at least k for any firm playing ν if at least one other firm undercuts and plays $\nu - \epsilon$. It holds that $r := \Pr(D_i \geq k|\nu; p_{min} < \nu) = \Pr((\frac{1-\lambda}{n})X \geq k) = 1 - \Pr(X < \frac{nk}{1-\lambda})$. Again using the distribution function of uniformly distributed random variables, one gets

$$r = 1 - \frac{nk}{2(1-\lambda)} \tag{3.6}$$

if $0 < k < 2\left(\frac{1-\lambda}{n}\right)$, as assumed. It can clearly be seen that r must be smaller than q for $\lambda > 0$, as should be expected.

As mentioned above, $2\left(\frac{1-\lambda}{n}\right)$ is the maximum demand (demand if $X = 2$) a colluding firm can get if some other firm deviates, as the fraction λ of informed consumers would never be attracted. Threshold values of k greater or equal to $2\left(\frac{1-\lambda}{n}\right)$ would thus result in a range of demand where collusion would be broken as punishment although firms know that every firm has colluded in the previous period with certainty. These paradox values of k are excluded in this analysis.

Next, since there is an expected total market demand of 1, the expected demand for a firm if every firm colludes is equal to $\frac{1}{n}$. This means that the expected profit during collusion is $\frac{\nu - c_L}{n}$ for every firm if cost is low. It was also calculated that the expected profit for a firm playing any price in the pricing support $[\underline{p}, \nu]$ must be equal to $(\nu - c_L)\frac{1-\lambda}{n}$ if there is no collusion. Once collusion ends and firms begin to price according to $F(p)$, each firm's profit will therefore be this high.

The only profit that is yet missing is the expected profit a firm makes when it deviates from a collusive strategy combination. The best possible deviation is to price at $p = \nu - \epsilon$, with ϵ close to zero. Doing so, a firm attracts all of the informed consumers without effectively facing a decrease of its profit margin. That is, a deviating firm will make an expected profit of $(\nu - c_L)\frac{\lambda n+1-\lambda}{n}$.

Overall, collusion under low cost can thus be maintained if the expected profit stream of a firm playing the strategy outlined on page 20 exceeds the expected profit stream of optimally deviating from the proposed strategy.[13] In any period, such an optimal deviation is characterized by playing $p = \nu - \epsilon$. Also, if deviation in any period is optimal, it must be optimal in every period since firms face the same maximization problem. This leads to the following

Proposition 3.1. *In the dynamic game, there exists a continuum of subgame perfect Nash-equilibria with demand threshold k if $\rho > \frac{n-1}{\lambda n}$. The profit maximizing threshold value k^* that is still incentive compatible with collusion is given by* $k^* = \frac{2(1-\frac{1}{n})(\frac{1}{\rho}-1)}{\frac{1}{1-\lambda}-n}$.

Proof. First, as firms' actions are restricted to pricing below or equal to ν, setting $p = \nu$ in every period of high costs is the only possible action firms can undertake under high costs. Pricing at ν under $c = c_H$ is thus clearly subgame perfect. Second, if costs are low, no matter what a firm has done in previous periods, it knows whether all of its competitors will keep pricing at ν in the current period (if demand has been greater than k in every elapsed period of low costs) or will price according to $F(p)$ (if demand has been low enough in some bygone low cost period).

[13] Recall that I assume for simplicity that firms do not discount.

In the former case, *if* colluding (i.e., pricing at ν) is incentive compatible for a firm given that all other firms price at ν, it must be incentive compatible no matter how many periods of low costs have already elapsed, as the random process costs follow is *memoryless*: for $\tau > 0$, it holds that $E(c^t = c_L|c^1, ..., c^{t-1} = c_L) = E(c^{t+\tau} = c_L|c^1, ..., c^{t+\tau-1} = c_L) = \rho$. This implies that firms find it optimal to keep colluding in *any* period of low costs *if* colluding is incentive compatible at all.
In the latter case, if a firm knows that every competitor will price according to the stage game equilibrium $F(p)$ until costs rise to c_H again, any price in the support $[\underline{p}, \overline{p} = \nu]$ constitutes a best response, as each price would yield the same expected profit. In particular, randomizing according to $F(p)$ must clearly be optimal.
I have thus proven that a trigger sales strategy as proposed will constitute a subgame perfect Nash-equilibrium of the dynamic game *if* colluding is incentive compatible for firms. When is that?
For collusion to be sustainable, it must hold that that the expected profit stream of collusion exceeds the expected profit stream of optimally deviating. Using the above results, it is clear that colluding in the first period where costs are low, given that all other firms collude, yields an expected profit of $\frac{\nu-c_L}{n}$, whereas deviating in the first period yields an expected profit of $(\nu - c_L)\frac{\lambda n+1-\lambda}{n}$.

Costs remain low for another period with probability ρ (with probability $1-\rho$, costs change to c_H and the firms' strategies are reset). Also, demand has only been higher than the critical threshold value k with probability q in the first period if every firm colluded. Thus, with probability q, a firm will continue to make an expected profit of $\frac{\nu-c_L}{n}$, whereas with probability $1-q$, it will only make the expected profit of the stage game equilibrium, i.e., $(\nu-c_L)\frac{1-\lambda}{n}$. In sum, the expected second period profit for a colluding firm is equal to $\rho[\frac{\nu-c_L}{n}q + (\nu - c_L)\frac{1-\lambda}{n}(1-q)]$.
On the other hand, if a firm deviated in the first period, demand only exceeded the

threshold value k with probability $r < q$. Thus, such a firm will continue to make the deviating profit of $(\nu - c_L)\frac{\lambda n + 1 - \lambda}{n}$ with probability r. With probability $1 - r$, a deviating firm will only make the expected profit of the stage game equilibrium $(\nu - c_L)\frac{1-\lambda}{n}$, as does the colluding firm with probability $1 - q < 1 - r$.

After this, a third period of low costs only happens with the *ex ante* probability of ρ^2. Also, since the measure of (potential) consumers is assumed to be *i.i.d.* across periods, it is obvious that the probability of firms having witnessed a demand greater or equal than k in both elapsed periods (of low costs) must be either q^2 (if the firm colludes) or r^2 (if the firm deviates), and so on.

Formally, for collusion to be sustainable it must therefore hold that

$$\begin{aligned}
\frac{\nu - c_L}{n} &+ \rho\left[\frac{\nu - c_L}{n}q + (\nu - c_L)\frac{1-\lambda}{n}(1-q)\right] + \\
&+ \rho^2\left[\frac{\nu - c_L}{n}q^2 + (\nu - c_L)\frac{1-\lambda}{n}(1-q^2)\right] + ... \geq \\
(\nu - c_L)\frac{\lambda n + 1 - \lambda}{n} &+ \rho\left[(\nu - c_L)\frac{\lambda n + 1 - \lambda}{n}r + (\nu - c_L)\frac{1-\lambda}{n}(1-r)\right] + \\
&+ \rho^2\left[(\nu - c_L)\frac{\lambda n + 1 - \lambda}{n}r^2 + (\nu - c_L)\frac{1-\lambda}{n}(1-r^2)\right] + ...
\end{aligned}$$

Multiplying by $\frac{n}{\nu - c_L}$ and using that $q^\tau + (1-\lambda)(1-q^\tau) = 1 - \lambda + \lambda q^\tau$, as well as $(\lambda n + 1 - \lambda)r^\tau + (1-\lambda)(1-r^\tau) = 1 - \lambda + \lambda n r^\tau$, this expression simplifies to

$$\begin{aligned}
1 &+ \rho(1 - \lambda + \lambda q) + \rho^2(1 - \lambda + \lambda q^2) + ... \geq \\
(\lambda n + 1 - \lambda) &+ \rho(1 - \lambda + \lambda n r) + \rho^2(1 - \lambda + \lambda n r^2) + ...
\end{aligned}$$

$$\begin{aligned} 1+\lambda[\rho q+(\rho q)^2+...] &\geq \lambda n+1-\lambda+\lambda n[\rho r+(\rho r)^2+...] \\ \lambda[1+\rho q+(\rho q)^2+...] &\geq \lambda n[1+\rho r+(\rho r)^2+...] \end{aligned}$$

which finally results in

$$\rho \geq \frac{n-1}{nq-r}. \tag{3.7}$$

Now, depending on the model parameters n and λ, two cases can be discerned when solving inequality (3.7) for k.

First, if $\lambda \leq 1-\frac{1}{n}$, no value of k exists that results in a trigger sales equilibrium for any $\rho < 1$. In consequence, the reversal of this inequality,

$$\lambda > 1-\frac{1}{n}, \tag{3.8}$$

gives a necessary condition for a collusive equilibrium under a trigger sales strategy. In case the above condition is not satisfied, the equilibrium of the stage game (i.e., all firms price according to $F(p)$) will be the only possible outcome in each low cost period *if* one restricts trigger sales strategies as the only alternative to $F(p)$.

Second, if $\lambda > 1-\frac{1}{n}$, it follows that

$$k \geq \frac{2(1-\frac{1}{n})(\frac{1}{\rho}-1)}{\frac{1}{1-\lambda}-n}. \tag{3.9}$$

Using that k is not allowed to reach or exceed $2\left(\frac{1-\lambda}{n}\right)$, the minimum ρ that is necessary for collusion can easily be computed. Setting $k = 2\left(\frac{1-\lambda}{n}\right)$ and solving condition (3.9) for ρ leads to

$$\rho > \frac{n-1}{\lambda n} := \rho_{min}, \tag{3.10}$$

which is the value that was stated in the proposition.

What is now the minimum value of k that allows for a collusive equilibrium? This value, say k^*, is given by inequality 3.9 solved for strict equality if (and only if) ρ exceeds the ρ_{min} that is given by condition (3.10). Formally, it holds that

$$k^* = \frac{2(1-\frac{1}{n})(\frac{1}{\rho}-1)}{\frac{1}{1-\lambda}-n} \text{ if } \rho > \frac{n-1}{\lambda n} \tag{3.11}$$

which is again the expression found in the above proposition.

If $\rho > \frac{n-1}{\lambda n}$, every trigger sales strategy that uses a k^* as defined in Equation (3.11) must constitute the optimal (in terms of firms' profits) collusive trigger sales strategy. I have thus proven that, given that the exogenous parameters n, λ and ρ of a market fulfill a certain criterium, a continuum of collusive, subgame perfect Nash-equilibria of the game exist. The optimal demand threshold k^* is given by the minimum value of k that is still incentive compatible with collusion.

□

A graphical example for parameter values that constitute collusive equilibria and the minimum collusive demand threshold k^* can be found in Figure 3.1.

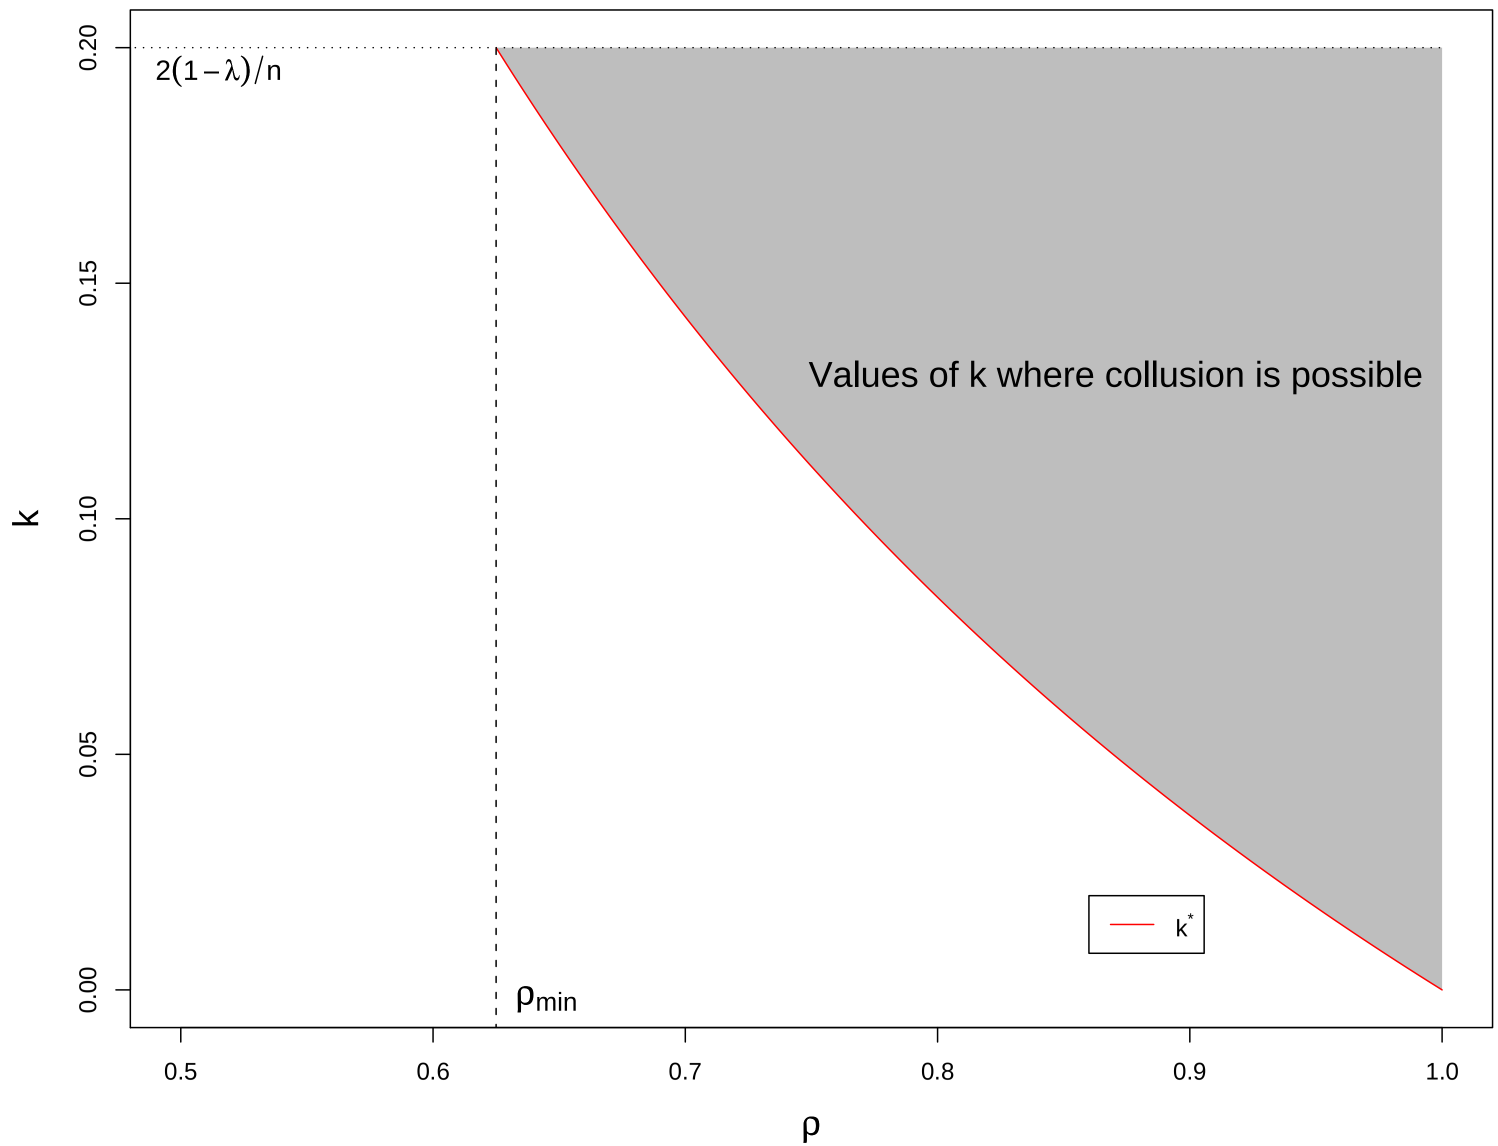

Figure 3.1 Values of k that constitute a trigger sales equilibrium for $n = 2$, $\lambda = 0.8$. Source: Own.

As a side note, reexpressing inequality (3.10) to $\rho > \frac{1-\frac{1}{n}}{\lambda}$, one can see that, given some fixed persistence parameter ρ, the fraction of informed consumers λ has to *increase* for an increasing number of firms in the market in order for collusion to be sustainable. This result is obtained because the absolute difference in profits under collusion (expected profit from collusion when some firm deviates vs. expected profit from collusion when all other firms collude) is given by $\frac{\lambda}{n}$ in each period, which is increasing in λ and decreasing in n. Given some fixed persistence parameter ρ, λ must therefore increase for an increased n in order to keep a colluding firm's ability to detect deviation on the same level.
The interpretation of this is that the signal extraction problem firms have to face when determining whether demand is low because of some random demand shock or because of a rival firm undercutting gets worse the smaller the number of informed consumers λ relative to the number of firms in the market n. If λ is small, firms have essentially no chance to successfully punish deviating firms and no equilibrium using a trigger sales strategy as proposed can exist. This signal extraction difficulty implies that a larger number of firms in the market needs to be balanced by a higher proportion of informed consumers.

Having determined k^*, in what follows, I will analyze the effect of the parameters ρ, n and λ on k^*. Then, I will translate the variable k^* back into a probability q^* that can be interpreted as the minimum probability with which firms end collusion because of random demand shocks in a cooperative equilibrium. This probability will finally be used to determine the speed with which negative cost shocks are passed along to prices.
The first thing one can see is that k^* is inversely related to ρ. The higher the persistence of costs, the lower minimum demand is needed to sustain collusion. This result is not surprising: the higher the probability that low costs persist, the higher the incentive for firms to keep colluding and not inducing a high risk of collusion to break by deviating.

Next, it is again straightforward to see that k^* is inversely related to λ. This shows that more informed consumers lead to a lower minimum demand needed for a collusive equilibrium. This property is consistent with the fact that more informed consumers make it easier for firms to distinguish between random demand shocks and demand shocks caused by a rival firm undercutting.
Finally, as the nominator of k^* is increasing and its denominator is decreasing in n, k^* is positively related to n. This means that the absolute demand threshold above which collusion can be maintained must *increase* in n although the expected number of consumers per firm under collusion decreases for a larger number of firms in the market. More competition thus leads to a smaller range of threshold values which are compatible with collusion, i.e., collusion gets more difficult until no collusive trigger sales equilibrium exists anymore. Similar results are typical for many models of industrial organization.

In order to understand the degree of asymmetric pricing in the model, I will now express the probability q^* of an additional collusive period that is associated with k^*. If k^* exists (i.e., Equation (3.11) holds), one can use Equation (3.5) to solve for q^*. Doing so, it follows that

$$q^* = \frac{\lambda - \frac{1}{\rho}(n-1)(1-\lambda)}{1-n(1-\lambda)}. \tag{3.12}$$

This leads to the following

Corollary 3.2. *If a collusive trigger sales equilibrium with threshold value k^* exists, the probability that collusion breaks down after at most j periods of low costs is given by $\psi(j) := 1-(q^*)^{j-1}$. This probability is*

- *decreasing with the persistence of costs ρ,*
- *decreasing with the fraction of informed consumers λ*

- *and increasing with the number of firms n*

in the market.

Proof. If every firm colludes at k^*, the probability that the random measure of consumers X falls below k^* will be the same in each period. Thus, considering that collusion can never break down in the first period of low costs, the probability that collusion is maintained for at least j periods is equal to $(q^*)^{j-1}$. Consequently, the probability that collusion breaks down after at most j periods must be given by $1-(q^*)^{j-1}$. The second part of the above corollary follows from $q = 1 - \frac{k}{2}$. Thus, the comparative statics of q^* must be opposite to those of k^*: the probability q^* that collusion is maintained for another period is increasing with ρ and λ and decreasing with n. As ψ is decreasing in q^*, ψ must finally be decreasing with ρ and λ and increasing with n.

□

Finally, it follows

Corollary 3.3. *Asymmetric pricing does emerge in the model as long as $\rho > \frac{n-1}{\lambda n}$. Positive cost shocks are passed along immediately to output prices whereas negative cost shocks are transmitted after an expected $\phi := \frac{1}{1-q^*} > 1$ periods.*

Proof. The mechanism that leads to asymmetric pricing is very straightforward. If a positive cost shock happens, margins are squeezed and firms have to increase their selling prices immediately to ν in order to avoid incurring losses (in fact, firms *have* to price at ν because their action space is restricted to pricing equal to or below ν).

If a negative cost shock happens, a collusive equilibrium will be played if $\rho > \frac{n-1}{\lambda n}$. Then, costs are sticky in the first period, meaning that no firm has an incentive to price lower than ν under collusion. In every period that follows where costs remain low, demand will be high enough to maintain collusion with probability q^*. With probability $1-q^*$, due to a random demand shock, demand drops sufficiently to lead firms to deviate as punishment, as they confuse random demand shocks with demand shocks caused by a rival firm undercutting. This implies that the probability that collusion breaks down exactly j periods *after* a negative cost shock is geometrically distributed with probability parameter $1-q^*$, i.e., given by $(q^*)^{j-1}(1-q^*)$ for $j > 0$ and zero if $j = 0$. As the expectation of a geometrically distributed random variable with probability parameter $1-q^*$ is given by $\frac{1}{1-q^*}$, this proves that prices will need an expected $\phi = \frac{1}{1-q^*} > 1$ periods to adjust after a negative cost shock, implying asymmetric pricing.

□

For example, for $n = 2$, $\lambda = 0.8$ (like in Figure 3.1) and $\rho = 0.75$, one obtains $q^* = 8/9$ and $\phi = 9$. Collusion will thus break an expected nine periods (e.g. days) after a negative cost shock, whereas prices will immediately jump to ν after a positive cost shock (if firms priced according to $F(p)$). A graphical depiction of the resulting probability *distribution* function of the number of low cost periods until collusion breaks down in the above example can be found in Figure 3.2.

While collusion will not break down deterministically in a collusive equilibrium, *once* collusion breaks down, every firm in the market will immediately price like in the competitive equilibrium. This would lead to an abrupt decline of average market prices once collusion ends, as can clearly be seen in the simulation I depict in Figure 3.3. In the next subsection, I will discuss the pattern of asymmetric price transmission that emerges when one market consists of multiple independent submarkets.

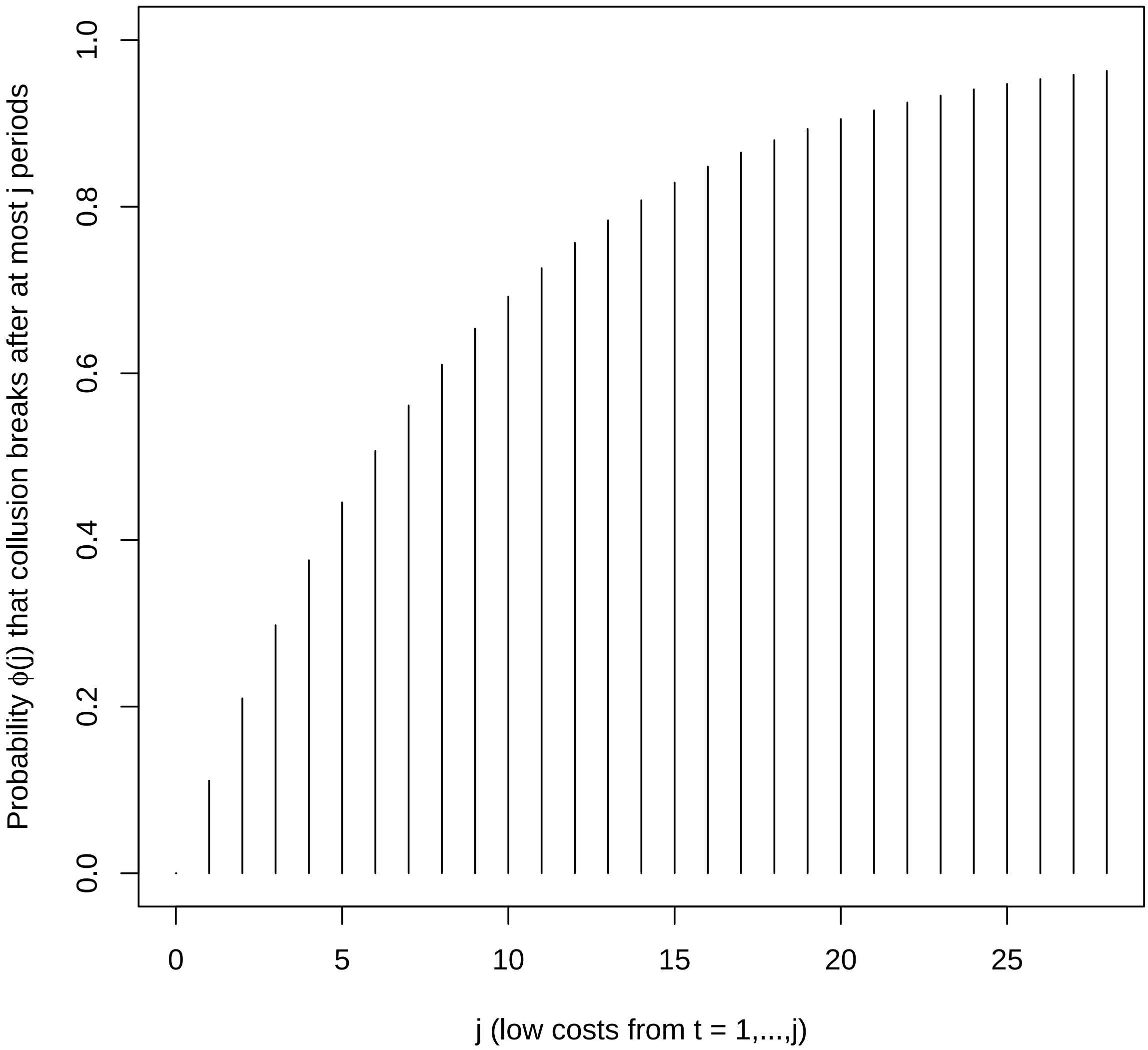

Figure 3.2 Probability distribution function of the number of periods until collusion collapses for $n = 2$, $\lambda = 0.8$, $\rho = 0.75$. Source: Own calculations.

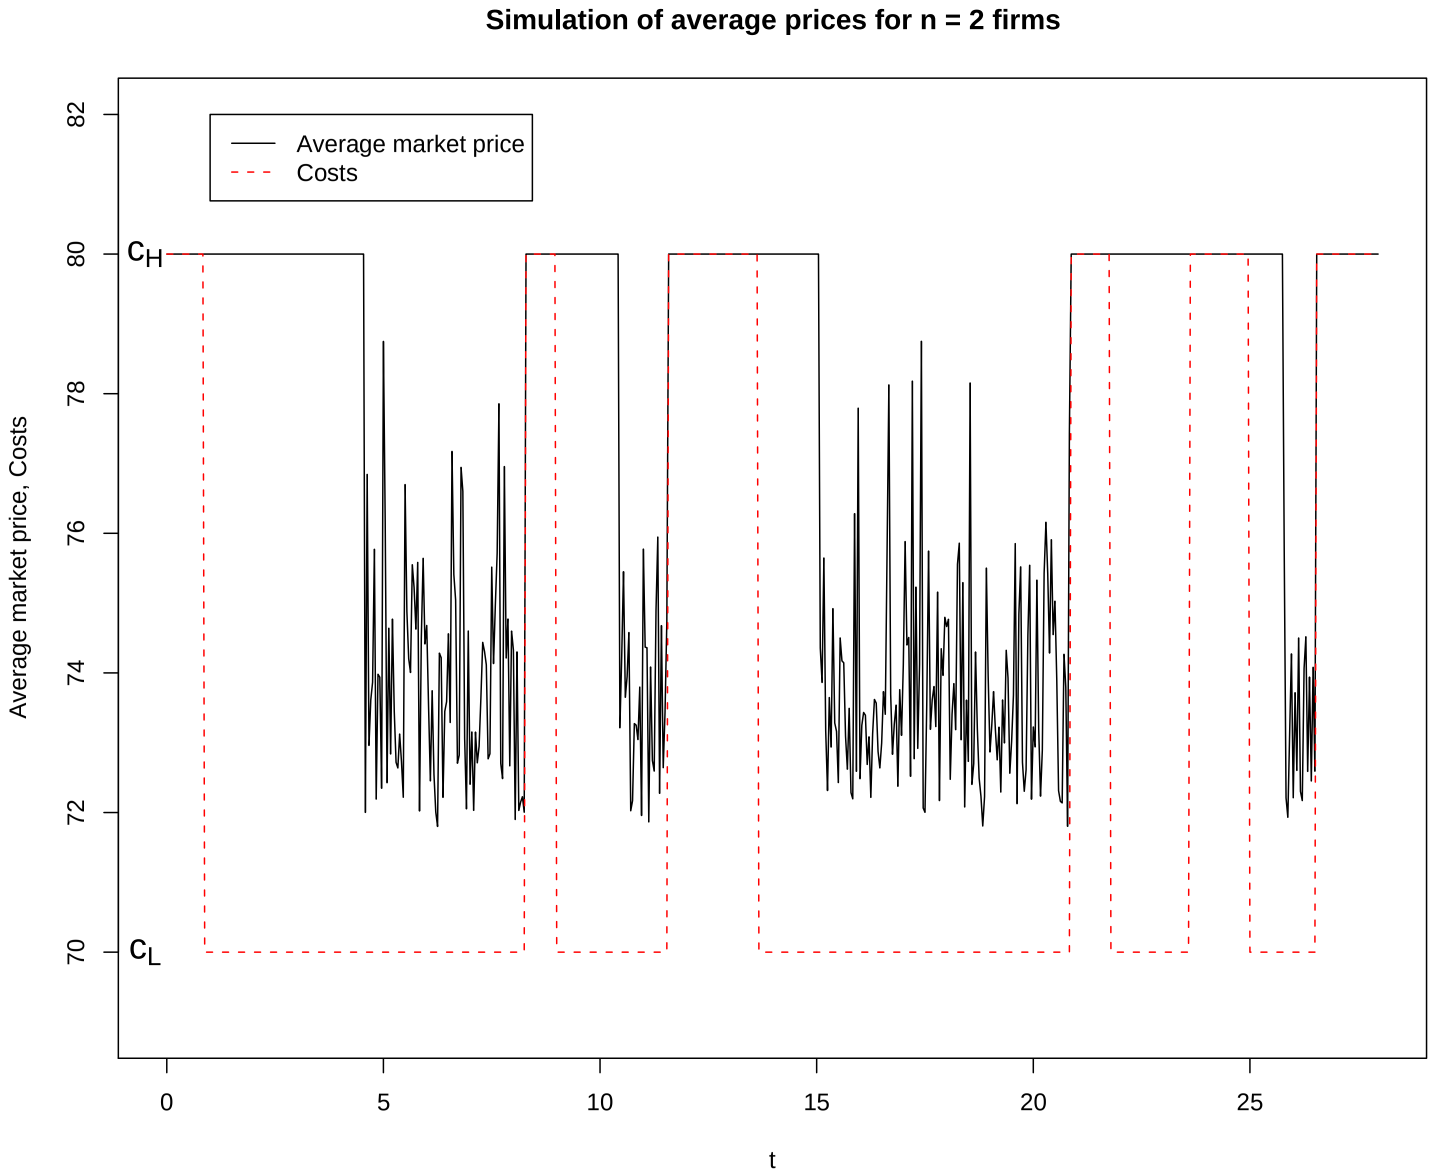

Figure 3.3 A simulation of average market prices for $n = 2$, $\lambda = 0.7$, $\rho = 0.98$. The interval between two periods can be interpreted as one hour, with a total of 28 days simulated. Asymmetric pricing is clearly apparent: prices will only adjust (in the sense that both firms in the market start pricing according to $F(p)$ instead of charging $p = \nu$) an expected 2.72 days after a negative cost shock, whereas positive cost shocks are transmitted without delay. Source: Own calculations.

3.4 Asymmetric Pricing with Separated Submarkets

For this subsection, consider the following extension of the model. Now, the *whole market* consists of a total of N firms who share the same cost level $c \in \{c_L, c_H = \nu\}$ in every period, with costs evolving according to the same two-state Markov process as in the basic model. The market is divided into $a = 1, ..., m$ spatially separated and completely independent submarkets with $n_a \geq 2$ firms each, such that $N = \sum_{a=1}^{m} n_a$. Firms compete through Bertrand competition in the submarkets only.
Also, for the consumer side, each market consists of $\lambda_a \in (0,1)$ informed consumers who observe all prices in the market, and $1 - \lambda_a$ uninformed consumers who only observe one price at random. Demand in each submarket follows the same principles outlined in Subsection 3.1.

In other words, each submarket is defined by the very same structure as the whole market in the basic model, and there is no interaction between the submarkets whatsoever. Because of this, the analysis of the extension of the basic model will be very simple: depending on the model parameters n_a and λ_a of a submarket, each result of the simple model will directly translate to the extended one: if $\rho \leq \frac{n_a - 1}{\lambda_a n_a}$, a submarket will be characterized by fierce competition and symmetric pricing, i.e., with every firm pricing according to the equilibrium of the stage game in each period. If, however, it holds that $\rho > \frac{n_a - 1}{\lambda_a n_a}$, a submarket will be characterized by a collusive trigger sales equilibrium with threshold value $k_a^*(n_a, \lambda_a)$ and asymmetric pricing, i.e., with firms adjusting prices to negative cost shocks after an expected $\frac{1}{1 - q_a^*(n_a, \lambda_a)}$ periods, in contrast to an immediate response to positive cost shocks.
Overall, the response of *average* prices to *negative* cost shocks will be different than in the simple model. Depending on how the parameters n_a and λ_a are distributed across submarkets, $0 \leq l \leq m$ of the submarkets will be characterized by collusive

equilibria and asymmetric pricing. On the other hand, $m - l$ of the submarkets will engage in fierce competition (where every firm plays $F(p)$ under low costs) and symmetric pricing.

As a consequence, the probability that $0 \leq \kappa \leq l$ of the l collusive submarkets end up in fierce competition after at most j periods of low costs would be complicated to determine, as different values of n_a and λ_a across submarkets would imply different probabilities $q_a^*(n_a, \lambda_a)$ with which collusion is maintained, which in turn would imply different values of the probability that collusion is maintained for at most j periods $\psi(j) = \psi(j; n_a, \lambda_a)$ for each submarket. Hence the number of collusive submarkets where collusion has ended after j periods would be a sum of l Bernoulli-distributed random variables with *different* probability parameter $\psi(j)$. While I will use values of n_a and λ_a that vary across submarkets in the simulation at the end of this subsection, I will not derive the probability with which collusion breaks down in κ of the l collusive submarkets after j periods of low costs if the collusive submarkets have different parameters n_a and λ_a.

However, if the parameters n_a and λ_a are the same for each collusive submarket, also q^* and hence $\psi(j)$ are the same for each submarket and the number of submarkets where collusion has broken down after (at most) j periods is simply binomially distributed with probability parameter $\psi(j)$. That is, it must hold that

$$\Pr(Z(j) = \kappa) = \binom{l}{\kappa} \psi(j)^{\kappa}(1 - \psi(j))^{l-\kappa} = \binom{l}{\kappa} [1 - (q^*)^{j-1}]^{\kappa}[(q^*)^{j-1}]^{l-\kappa} \quad (3.13)$$

with $Z(j)$ denoting the number of submarkets where collusion has broken down after j periods of low costs.

Also, in expectation, $l\psi = l[1 - (q^*)^{j-1}]$ of the collusive submarkets will engage in fierce competition after j periods, which is increasing in j.

Figure 3.4 depicts the probability mass function of the number of collusive submarkets where collusion has collapsed after $j = 2, 3, 5, 10$ and 20 periods of low costs. The parameters involved are: $n_a = 4$ and $\lambda_a = 0.9$ for every collusive submarket (implying that $q^* = 0.782$ for every submarket), $\rho = 0.85$, $l = 50$. One can clearly observe that the expected number of firms that quit collusion increases for the number of elapsed low cost periods j. As the number of submarkets where collusion is maintained will reduce gradually in expectation, *average* market prices for the whole market will decrease in expectation if $c = c_L$ persists.

Thus, given that there are l collusive and $m - l$ fierce submarkets, a negative cost shock will have the following effect:

- In the first period, prices will drop to $F(p)$ for all the submarkets which are characterized by fierce competition, implying that average market prices drop to some extent in the first period of low costs.
- In every period that follows, some of the collusive submarkets will quit cooperation in expectation, leading to a gradual decline of average market prices.
- This goes on until either every firm in the whole market prices according to $F(p)$ (implying that average market prices have reached their minimum) or costs rise back to c_H.

Finally, the pattern of asymmetric price transmission that is observed in many markets becomes apparent: Positive costs shocks are immediately passed along to output prices whereas negative cost shocks are only partly passed along to output prices in the first period and then continue to decline gradually in expectation until their minimum is reached when every firm prices according to $F(p)$.

Figure 3.5 shows a simulation of the behavior of average market prices during a period of four weeks (28 days), in which each hour a cost shock can happen with

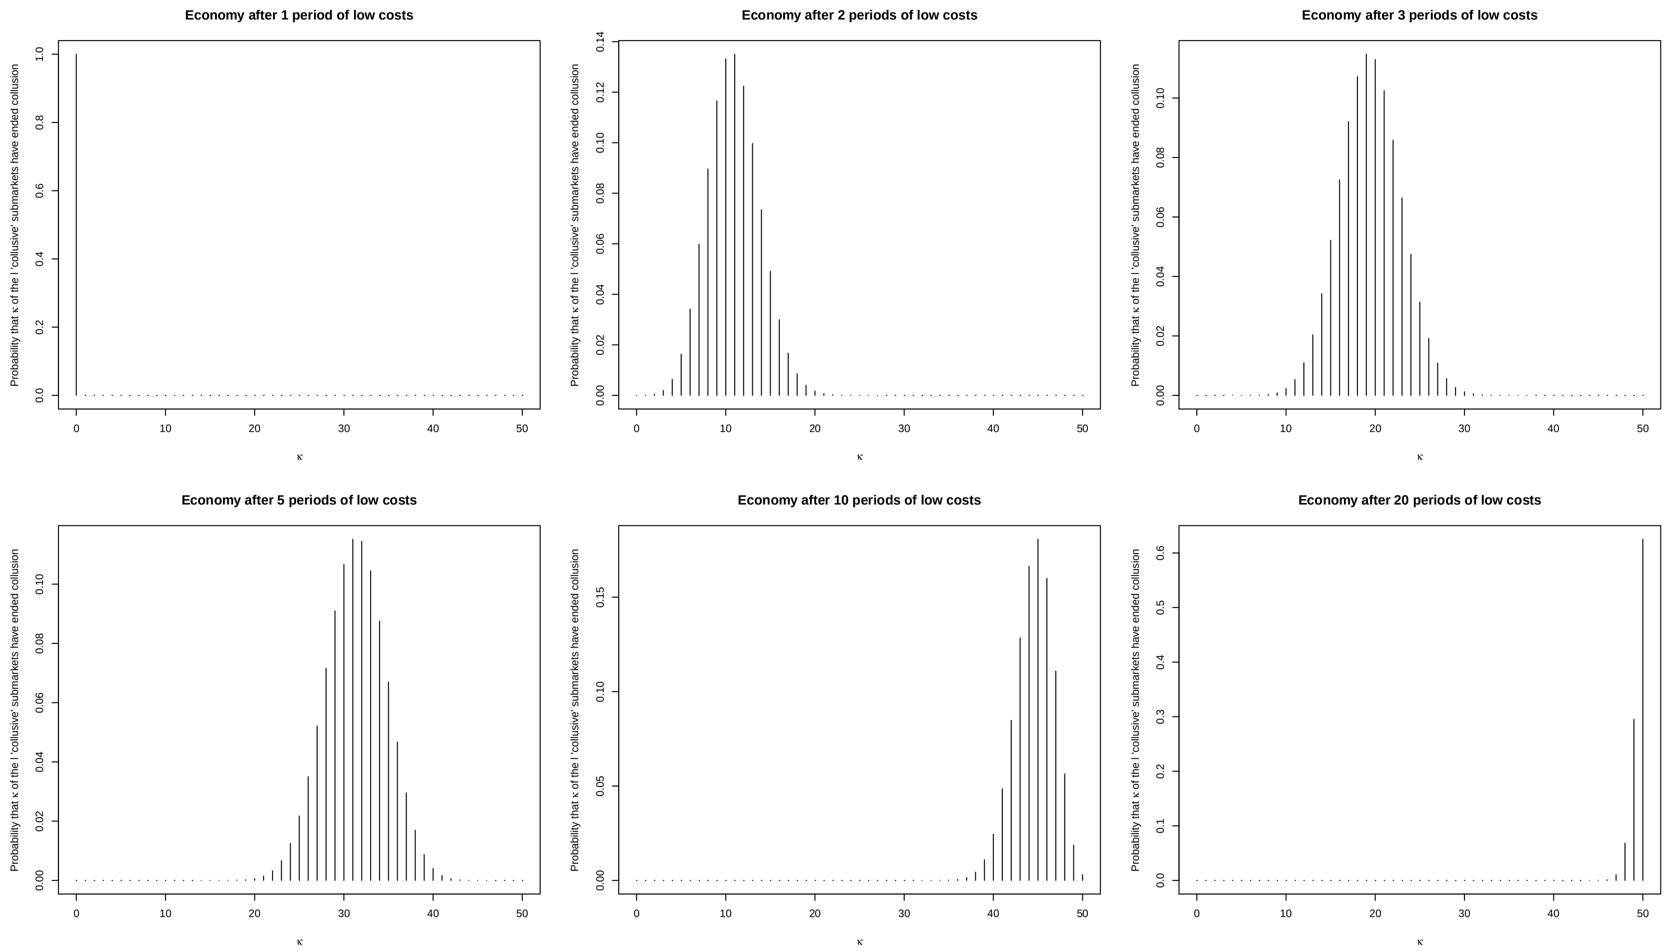

Figure 3.4 Probability mass function of the number of submarkets κ who have broken collusion after $j = 2, 3, 5, 10$ and 20 periods for $n_a = 4$ and $\lambda_a = 0.9$ in every submarket, $\rho = 0.85$, $l = 50$. Source: Own calculations.

probability 0.02. The following parameters were used: $c_L = 70$, $c_H = \nu = 80$, $\rho = 0.98$, $n_a \sim P(1) + 2$ (implying an expected three firms in every submarket), $\lambda_a \sim U[0.6, 1]$ and $m = 50$. The overall number of firms that was the result of this simulation was $N = 151$, while $l = 41$ of the $m = 50$ submarkets were collusive.

One can see that prices do indeed follow an asymmetric transmission pattern: average prices rise instantaneously to ν if costs rise to c_H. In contrast, after negative cost shocks, there is only a partial immediate effect on average prices because out of the $m = 50$ submarkets in the market, only nine are characterized by fierce competition (with those immediately pricing according to $F(p)$ whenever costs drop to c_L).
After this initial reaction, average prices continue to decline because of random demand shocks that lead collusion to successively break down in the other, collusive submarkets. Overall, asymmetric pricing is present: negative cost shocks need much longer to be fully[14] transmitted to output prices than positive ones.

The corresponding pattern of how collusion breaks down in the submarkets in order to produce Figure 3.5 can be found in Figure 3.6.

[14]It has to be noted that even in the competitive scenario, average prices will never reach costs under the low cost regime: after all, firms price according to $F(p)$ with a minimum price of $\underline{p} > c_L$ (see Subsection 3.2). Because of this, average prices will typically be somewhere in the middle of c_L and $c_H = \nu$ even if none of the firms collude.

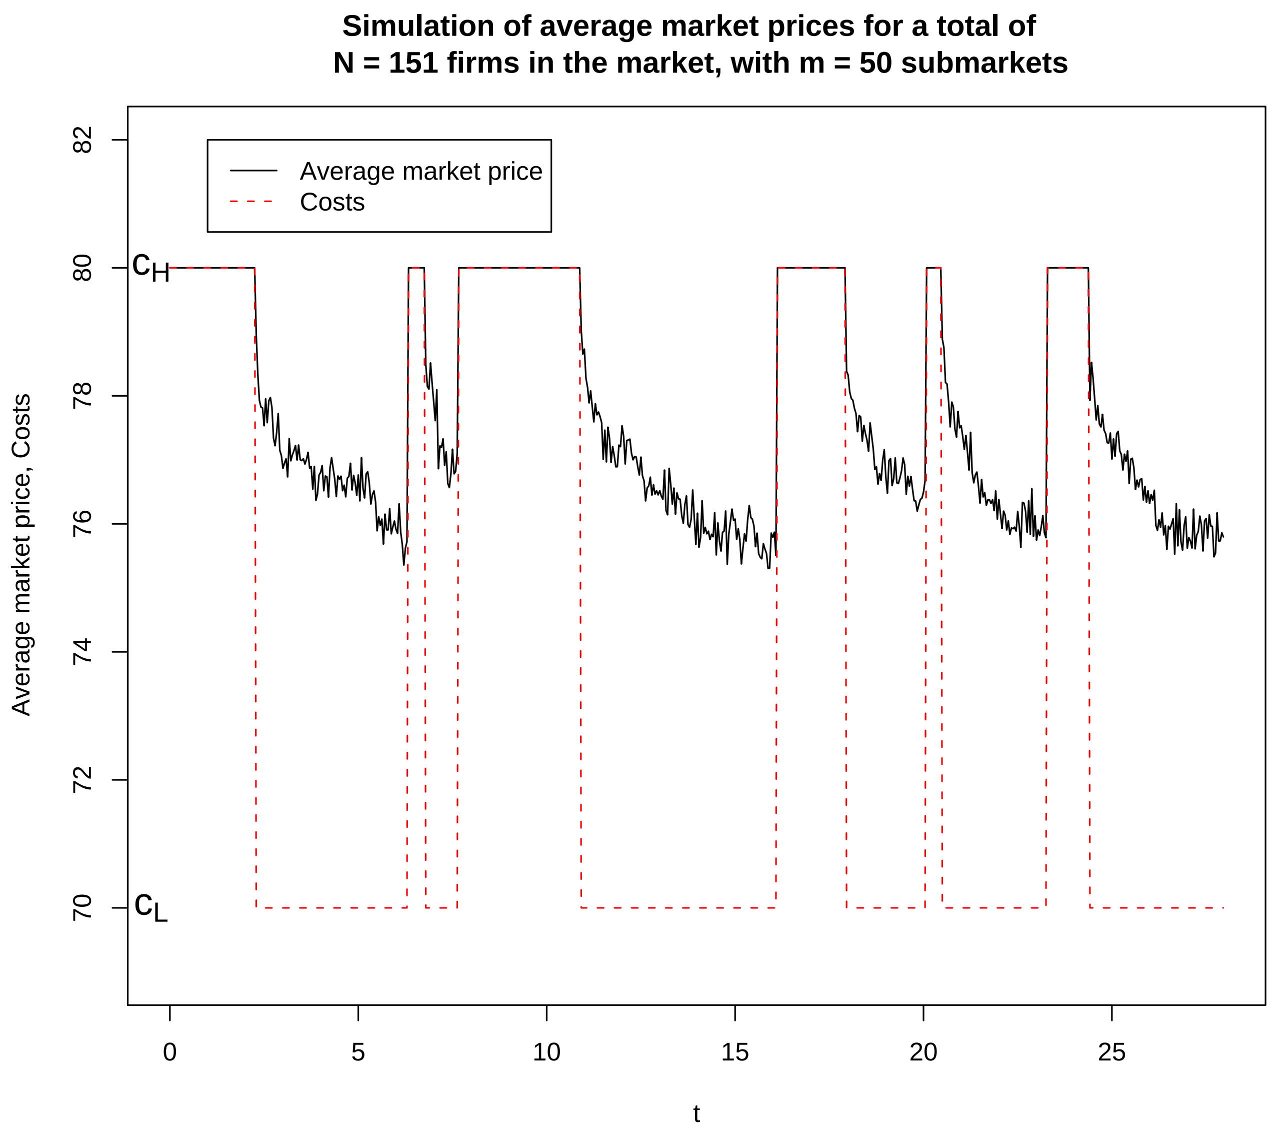

Figure 3.5 A simulation of average market prices for $c_L = 70$, $c_H = \nu = 80$, $\rho = 0.98$, $n_a \sim P(1) + 2$, $\lambda_a \sim U[0.6, 1]$ and $m = 50$.
Source: Own calculations.

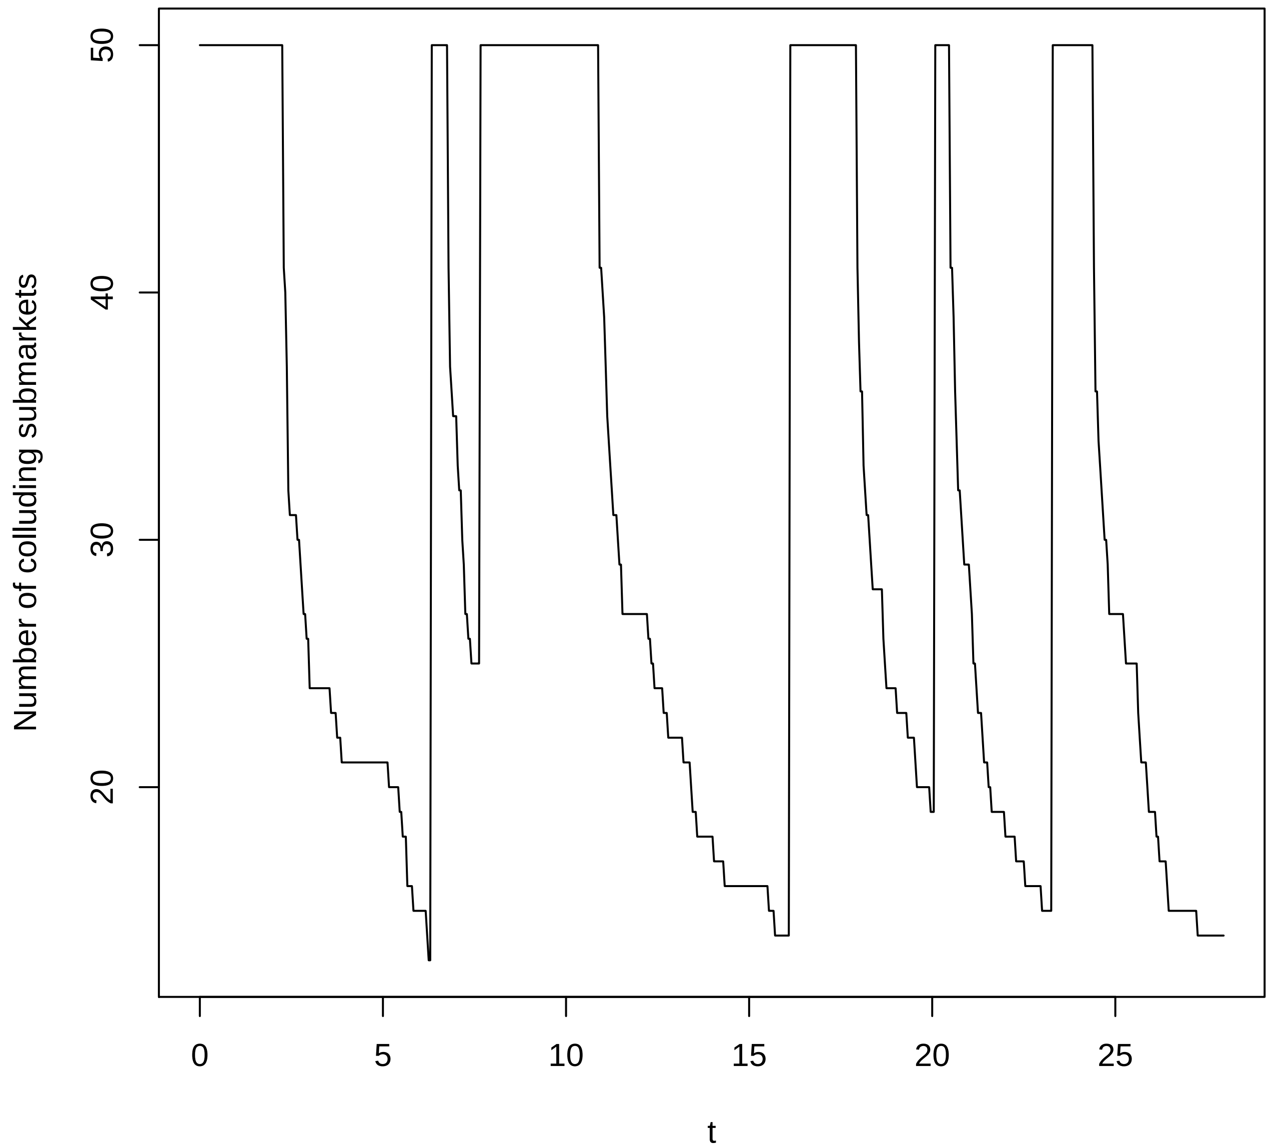

Figure 3.6 A simulation of how collusive submarkets collapse for $c_L = 70$, $c_H = \nu = 80$, $\rho = 0.98$, $n_a \sim P(1) + 2$, $\lambda_a \sim U[0.6, 1]$ and $m = 50$. Source: Own calculations.

3.5 Summary and Discussion

In this section, I have developed a trigger sales model able to generate asymmetric price transmission caused by oligopolistic coordination. After negative cost shocks, firms use the price charged under high costs as focal point for collusion. However, firms confuse random demand shocks with demand shocks caused by rival firms undercutting. In consequence, collusion can break down although every firm wishes to collude. In the simple model, collusion will break down at once, leading to a sudden and full decline of average market prices once collusion ends. In contrast, if one assumes that markets are composed of several spatially separated submarkets, average market prices can decline smoothly after a negative cost shock. In the first period after a shock, prices will drop to some extent because of submarkets where collusion is impossible. In every low cost period that follows, the expected number of submarkets where collusion has ended increases, implying that average market prices will decline in a smooth fashion if the number of submarkets in a market is relatively large.

There are several implications of my model that can be tested empirically. First, collusive equilibria and thus asymmetric pricing under collusion should only be found if the number of informed consumers is large, relative to the number of firms in a market. Even for small numbers of firms in a market, high fractions of informed consumers are needed to allow for asymmetric pricing. Next, I calculated that the speed with which negative cost shocks are transmitted to output prices decreases for a higher persistence of costs and larger fractions of informed consumers, but increases for a larger number of firms in a market.
While the first and third of these comparative statics are intuitively straightforward, the second result is somewhat strange: the more informed consumers there are in a market, the faster cost changes are generally expected to be reflected in

output prices. However, a larger fraction of informed consumers does in fact facilitate collusion in the model as it makes it easier for firms to monitor the action of their competitors. This effect outweighs the increased incentive for firms to deviate from a collusive equilibrium, which in turn leads a larger proportion of informed consumers to imply a slower response of output prices to negative cost shocks.

As a final note, I want to say that the purpose of this exercise was to generate a simple and straightforward model of asymmetric pricing under collusion. The main strength of my approach is that it provides a basic mechanism of rockets and feathers that future models can build on. Also, even the simple theoretical framework I used is able to generate several economic implications that can be tested empirically. By comparing these features with the predictions of other models of asymmetric pricing, researchers might be able to better distinguish between the causes of asymmetric pricing in different markets. This in turn could improve policy makers' efficiency to reduce the welfare redistribution from consumers to producers that is associated with asymmetric price adjustment.

Like in any model of real markets, the simplicity of my model comes at a cost. For example, the assumption that there are only two possible cost states, with high costs being equal to the reservation price of consumers, is quite unrealistic. In actual markets, costs can take a continuum of values, with cost shocks being randomly large. Also, the value of the reservation price ν will usually differ for individual consumers, implying that demand should be elastic to prices even if every firm colluded at the same price. However, I think that the basic intuition of my model should carry over if one extended it to a continuum of cost states.[15]

[15] A possible solution to the resulting coordination problem firms would have to face when costs drop might in fact be to assume demand that is price elastic, as this implies some unique monopoly price given costs.

More importantly, my model doesn't even try to explain why demand might be random across periods or why there is exactly a fraction λ of informed consumers in the market. For more realism, one would have to incorporate consumer search and thus, the search intensity of consumers (and possibly overall demand) would become endogenous.

Finally, I think that the model could be improved if one allowed for an interaction between the separated submarkets. If, for example, the measure of potential consumers in a submarket was somehow negatively related to the number of submarkets with fierce competition, different dynamics for the transmittance of negative cost shocks could be obtained. Then, prices might begin to adjust slowly, but as collusion collapses in more and more submarkets, the probability that some collusive submarket still maintains a demand greater than the minimum threshold diminishes, as more and more potential consumers leave a submarket.

To sum up, I think that my model might capture some of the effects that drive asymmetric pricing under oligopolistic coordination. While some of the model's assumptions are clearly unrealistic and numerous extensions of the model can be conceived, most of its implications seem plausible and (as can be seen below) are also shared by other models of asymmetric price transmission. Thus, I think that my model provides a good starting point for future research in the area of asymmetric pricing caused by oligopolistic coordination.

4 Asymmetric Price Adjustment due to Consumer Search

Although policy makers, the media and a fair share of researchers tend to attribute asymmetric pricing to an abuse of market power and oligopolistic coordination, it was already mentioned in the introduction of this thesis (see Section 1) that consumer search coupled with asymmetric information drives most contemporaneous models of rockets and feathers. More precisely, there are at least four formal models capable of generating asymmetric pricing through imperfect consumer search.[16]

In this section, I will try to give the reader insight into these theoretical approaches. Doing so, in Subsection 4.1 and Subsection 4.2, I will explain the models of Yang and Ye (2008) and Tappata (2009) in detail, both providing their rudimentary mathematical structure and discussing some of their unorthodox aspects as well as possible extensions. In both subsections covering these models, I will start to outline the basic mechanism that drives asymmetric pricing in the model, explain the model's setup and derive equilibrium of the static game. This will be followed by the derivation of equilibrium of the dynamic game and, for the model of Yang and Ye, the computation of several comparative statics. A short discussion concludes.

Finally, in Subsection 4.3, I will briefly provide the intuition behind the two other consumer search models of asymmetric pricing that are currently available. These were created by Lewis (2009) and Cabral and Fishman (2008). For the sake of brevity, I will not include any derivations in the discussion of these models.

[16]Not counting the dynamic oligopoly model I developed in Section 3, I am only aware of two other formal models that lead to asymmetric pricing. These two models (as well as the idea of Borenstein et al. (1997) that asymmetric pricing might be the product of finite inventories and lags in adjustment of production) will be covered in Section 5.

4.1 A Model of Consumer Search With Learning

Model Setup and Equilibrium of the Static Game The first consumer search model that leads to asymmetric pricing I will discuss in this thesis was suggested by Yang and Ye (2008). It is a dynamic model of search with learning. In contrast to the traditional opinion that asymmetric pricing is a consequence of market power (see Section 3), the model shows that there can be a natural tendency towards it even in markets where no collusion is apparent. This tendency is caused by asymmetric information between firms and consumers, with the latter being unable to directly observe the cost realization of firms.

The basic principle that drives asymmetric pricing in the model is that positive cost shocks are immediately learned by agents who search, which results in a full adjustment of the search intensity and prices in the next period. Conversely, non-searchers need longer to learn the true cost state when a negative cost shock occurs, which leads to a slower adaption of the search intensity and prices.

The model's setup is as follows. The agents in the model are a continuum of rational[17] consumers and a continuum of firms (having capacity constraints) producing a homogeneous good. All firms share the same unit cost level, which can be either high or low and is unobserved by consumers. There are three types of consumers: consumers who always search (low search costs), consumers who never search (high search costs), and, most importantly, *critical* consumers (intermediate search cost) that endogenously determine whether they search or not. Like in other search models, these consumers will search when the expected benefit of search (in terms of expected price reduction) exceeds their search costs.

[17]In the sense that they optimally process all the information available to them.

Nonsearchers will shop randomly and buy from the first firm they encounter (given there is no binding capacity constraint for the firm), while searchers will observe the price quotations of all firms and buy at the firm with the lowest price (again given that this firm is not capacity constrained).[18] That is, the protocol of *nonsequential* search is adopted.

Formally, the total measure of firms is normalized to one, with *each* firm having a (marginal) cost realization of either c_L or c_H, with $c_L < c_H$. Firms observe this realization at the beginning of the static game and then compete in prices. Also firms have a capacity constraint in the sense that each firm can sell $k < \infty$ units of the good at most.

The continuum of consumers is normalized to a total measure of β, with $\beta > 1$. β can thus be interpreted as the number of consumers per firm in the market. Consumers have a unit demand with a reservation price equal to $\nu > c_H$. It is also assumed that $\beta < k$, which implies that the number of consumers per firm in the market is smaller than the capacity constraint (i.e., overall, there are enough goods for each consumer). Consumers cannot observe the true cost realization of firms, but have certain beliefs. Let α denote the probability a consumer assigns to the realization of the high cost level, c_H.
The type of each consumer is fully characterized by his or her search cost. A fraction λ_1 of consumers has zero search cost $s_L = 0$; these consumers are named

[18]The capacity constraint of firms that is part of the model can result in rationing among consumers, i.e., it can happen that low-price firms attract too many customers. In this case, it is assumed that each customer is served with equal probability. If a nonsearcher is rationed, they will randomly shop at another firm (without additional cost). If a shopper is rationed, they will go to the firm with the lowest price of all remaining firms (and so on, if they are rationed again).

shoppers. Each shopper will search in equilibrium, no matter what their beliefs α are. A fraction λ_2 of consumers is assumed to have a search cost $s_H > \nu$, implying that it will never search in equilibrium. The remaining $1-\lambda_1-\lambda_2$ consumers have intermediate search cost $s_M \in (s_L, s_H)$. Their decision whether to search or not is endogenously determined and depends on their beliefs about the cost state α. As mentioned above, they are referred to as critical consumers.
As the critical consumers are the only consumers that have an endogenous search decision, I will denote the cumulative distribution function of critical consumers' beliefs by $F(\alpha)$, with $F(\alpha)$ being the fraction of critical consumers believing $\Pr(c = c_H) \leq \alpha$.

The stage game will have the following timeline. After observing production costs, firms simultaneously set their prices. Then, the critical consumers make their search decision based on their beliefs, while all of the shoppers search for sure and all the nonshoppers don't. The game is complete after the buying process of consumers is finished.

The focus of the analysis will be laid on symmetric equilibria, which means that each firm has to use the same pricing strategy, following a price distribution G. If one denotes the endogenously determined $\mu \geq \lambda_1$ as the proportion of consumers who search (shoppers plus critical consumers that search), a symmetric perfect Bayesian equilibrium of the game can be characterized by a combination of μ^* and $G^*(.|c)$ such that given the equilibrium search intensity μ^* and their cost level, firms' optimal pricing strategies yield the price distribution $G^*(.|c)$, while given $G^*(.|c)$ and the consumers' beliefs $F(\alpha)$, consumer's optimal search decisions lead to a search intensity of μ^*.

As a starting point, the analysis of the game is carried out with a fixed (exogenous) search intensity μ. Let $\underline{p}$ be the lowest price in the support of the equilibrium price distribution, and $\eta(\underline{p})$ be the *proportion* of firms charging this price. Then a firm charging $\underline{p}$ will make sales of

$$\frac{\mu\beta}{\eta(\underline{p})} + (1-\mu)\beta \tag{4.1}$$

if this number is smaller than it's capacity constraint k, or sales of k if not.[19]

The authors prove that there are several direct implications of the above result.[20] First, they show that firms charging $\underline{p}$ must always sell k units in every equilibrium, and that they are not rationed while doing so. Formally, it must thus hold that each low price firm gets exactly a demand of k in equilibrium, i.e., it holds that

$$\frac{\mu\beta}{\eta(\underline{p})} + (1-\mu)\beta = k. \tag{4.2}$$

[19]The most natural way to understand Equation (4.1) is by thinking in discrete terms. If the total *number* of consumers is denoted by n_c, the total number of firms is denoted by n_f and there are $n_f(\underline{p})$ firms charging a price of $\underline{p}$, the number of searchers shopping at some random firm that prices $\underline{p}$ would clearly be $\frac{\mu\, n_c}{n_f(\underline{p})}$.

By expanding the numerator and denominator of this ratio by $\frac{1}{n_f}$ and using that $\frac{n_c}{n_f} = \beta$ as well as $\frac{n_f(\underline{p})}{n_f} = \eta(\underline{p})$, one can see that

$$\mu\frac{\frac{n_c}{n_f}}{\frac{n_f(\underline{p})}{n_f}} = \frac{\mu\beta}{\eta(\underline{p})},$$

which is the first term of Equation (4.1).

Also, a firm pricing $\underline{p}$ can expect to attract $\frac{(1-\mu)n_c}{n_f}$ nonsearching consumers by chance. Using again that $\frac{n_c}{n_f} = \beta$, the second term of Equation (4.1) becomes clear.

[20]For a rigorous proof, see Yang and Ye (2008, p. 552f)

Also, it can be shown that no equilibrium exists such that prices are *continuously* distributed over $[\underline{p}, \overline{p}]$ if $\overline{p} \leq \nu$. From this the most important property can be derived, namely that

Lemma 1. *The price distribution has two mass points. Given μ and c, each firm must either charge $p = \nu$ or $p = \underline{p} \in (c, \nu)$ in equilibrium.*

Because firms can only choose between pricing at ν and pricing at $\underline{p}$, the final step to obtain an equilibrium of the static game when the search intensity μ is exogenous is to equate the profits of firms charging these two prices, as both types of firms have to make the same profit in equilibrium. As mentioned above, a low price firm will sell k units and thus make a profit of $\pi(\underline{p}) = k(\underline{p} - c)$. A high price firm will face a demand of $(1-\mu)\beta$ and hence make a profit of $\pi(\nu) = (1-\mu)\beta(\nu - c)$. Equilibrium is then defined by Equation (4.2) and

$$\pi(\nu) = \pi(\underline{p}). \tag{4.3}$$

Rearranging equations (4.2) and (4.3), it directly follows that

$$\eta(\underline{p}) = \frac{\mu\beta}{k - \beta + \mu\beta} \tag{4.4}$$

and

$$\underline{p} = c + \frac{(1-\mu)\beta}{k}(\nu - c). \tag{4.5}$$

It can easily be proven that the price distribution from above (a fraction $\eta(\underline{p})$ of the firms pricing $\underline{p}$ and the remaining firms pricing ν) is in fact an equilibrium of the static game, given an exogenously determined μ, some cost level c and the parameters β, k and ν. To do so, one can simply show that it makes no sense to deviate for any firm pricing ν or $\underline{p}$.

It is worth noting that all the properties of the above equilibrium make sense. $\eta(\underline{p})$ is an increasing function with μ, implying that a higher search intensity will lead more firms to price low. Also, the low price $\underline{p}$ is decreasing with μ. The interpretation for this is that when more consumers decide to search, the overall elasticity of demand will increase, making it profitable for some high price firms to deviate. Low price firms will (cet. par.) attract more consumers, implying that their profit margin $\underline{p} - c$ must become smaller (by a reduction of $\underline{p}$) to keep them indifferent between pricing at $\underline{p}$ and ν.
It can also be seen that the cost level c has no influence on the ratio of low and high price sellers, but only on the level of the low price.

Having determined the equilibrium of the static game when the search intensity is exogenous, it is finally necessary to endogenize this variable in order to derive a true equilibrium of the stage game. For this, it is assumed that firms know the distribution of consumers' beliefs $F(\alpha)$, while consumers may or may not know it. Then, it is possible to determine the expected price reduction for a consumer with belief α, i.e., who beliefs that $\Pr(c = c_H) = \alpha$ and $\Pr(c = c_L) = 1 - \alpha$. Doing so, the authors show that the expected price reduction by search is not influenced by the search intensity μ, which means that the gains from search are independent from the search level. It also becomes apparent that the expected gain from search is decreasing with a consumer's belief α. This is quite an intuitive result and shows that consumers with a pessimistic belief (i.e., high α) expect to have lower gains from search (compared to more optimistic consumers). It also follows that there must be a cutoff belief $\hat{\alpha}$ which divides the group of critical consumers into searchers and nonsearchers. $\hat{\alpha}$ can simply calculated by equating the search costs of the critical consumers s_M with their expected gains from search, which results in

$$\hat{\alpha} = \frac{(\nu - c_L)(k - \beta) - ks_M}{(c_H - c_L)(k - \beta)}. \tag{4.6}$$

From this, it is only a small step to calculate the equilibrium search intensity μ^*. As $F(\alpha)$ is defined as the fraction of consumers that have beliefs of α or less (with less being more optimistic), it follows that a fraction $F(\hat{\alpha})$ of the critical consumers will search. Combined with the fraction λ_1 of consumers that will search anyway, the equilibrium search intensity μ^* can be written as

$$\mu^* = \lambda_1 + (1 - \lambda_1 - \lambda_2)F(\hat{\alpha}). \tag{4.7}$$

The authors show that both $\hat{\alpha}$ and μ^* are unique, which gives rise to a unique equilibrium. This unique equilibrium of the static game is defined by equations (4.4), (4.5), (4.6) and (4.7). Given $F(\alpha)$ and the parameters of the model, consumers' optimal decisions lead to an equilibrium search intensity of μ^*. Firms anticipate that and price accordingly, following the distribution derived earlier. No rationing occurs in equilibrium.

To finish the analysis of the static game, it is useful to look at some attributes and comparative statics of the equilibrium. The first thing one can see is that the average price paid by customers will be lower under low costs. This is because the low price $\underline{p}_L$ under low costs (when $c = c_L$) is lower than the low price $\underline{p}_H$ under high costs (when $c = c_H$), while the high price ν and the fraction of low and high price firms is the same in both states. It can also be shown[21] that prices are more *dispersed* (as measured by the difference of the average price to the lowest price)

[21] see Yang and Ye (2008, p. 554)

under the low cost state, making search more rewarding under this state (although consumers don't know in which state they are, of course).

Also, as $F(\hat{\alpha})$ is increasing in $\hat{\alpha}$, a higher cutoff level for critical consumers $\hat{\alpha}$ will imply a higher search intensity μ^*. From this it follows that $\underline{p}_H$ and $\underline{p}_L$ are both decreasing in $F(\hat{\alpha})$ (as $\underline{p}$ is decreasing in μ). Furthermore $\eta(\underline{p})$ is increasing in $F(\hat{\alpha})$ too because $\eta(\underline{p})$ is increasing in μ. Finally it can be proven that the average price paid by consumers is decreasing in μ and increasing in $F(\hat{\alpha})$. This concludes the analysis of the static game.

Equilibrium of the Dynamic Game and Comparative Statics Now, the static game as described above will be extended to a dynamic framework. For this, I consider the static game played repeatedly over time. Time is modeled as discrete, i.e., $t = 1, 2, ...$, and will be denoted by a subscript.

The first thing necessary to model the dynamic game is to describe how costs will evolve over time. Like in the dynamic oligopoly model of Section 3, it is assumed that

$$\Pr(c_{t+1} = c_H | c_t = c_H) = Pr(c_{t+1} = c_L | c_t = c_L) = \rho > 0.5, \tag{4.8}$$

i.e., that there is some persistence parameter $\rho > 0.5$ that defines the probability of observing the same cost state in $t + 1$ as in t. For the first period ($t = 1$), it is assumed that $\Pr(c_1 = c_H) = \Pr(c_1 = c_L) = 0.5$, that is both states are equally likely at $t = 1$. Also, Equation (4.8) and the probabilities for the first period are considered common knowledge.

It is further assumed that firms always know the current and past cost levels whereas consumers can only update their *beliefs* based on the prices they observe. Consumers do not get to know the past cost realizations, although they are sometimes able to exactly determine them from the price distribution. The assumption that consumers do not eventually find out past cost realizations is a very essential one for the model. Only because of this, there can be a gradual learning process among consumers that will lead to asymmetric price adjustment. In the next section, a model of Tappata (2009) will be discussed that doesn't incorporate learning and hence explains asymmetric price adjustment in a different fashion.

Next, it is important to recapitulate that the price distribution of the static game was dependent on consumers' beliefs and the cost realization. In particular, one could see that $\underline{p}$ was dependent on the cost level c, and it could be shown that $\underline{p}_H > \underline{p}_L$ holds for any search intensity μ. The dynamic setting can now use this feature to explain a delayed response to negative input price changes.

To do this, it is helpful to make some parameter restrictions that greatly reduce the complexity of the problem first. The parameters are assumed to be such[22] that the lowest possible price under high costs, i.e., $\underline{p}_H$ when all critical consumers search, will still be higher than the highest possible price under low costs, i.e., $\underline{p}_L$ when none of the critical consumers search (recall that $\underline{p}$ is negatively related to μ). This can also be called the *nonoverlapping condition*, meaning that the supports of $\underline{p}_H$ and $\underline{p}_L$ (whose ranges are determined by μ) do not overlap.
The simple conclusion to this is that a price observation of $\underline{p}$ directly tells a consumer in which cost regime they are. Using Equation (4.8), such a consumer can immediately calculate the correct probabilities for the cost states in the next period.[23] In contrast, a (nonsearching) consumer that observes the high price ν will not be able to infer anything from this observation because ν is the high price in each cost regime and the fraction of firms charging ν, $1-\eta(\underline{p})$, doesn't depend on

[22]See Yang and Ye (2008, p. 555)

[23]Their belief $\alpha_{t+1} = \Pr(c_{t+1} = c_H)$ will be ρ if they observe $\underline{p}_H$ and $1-\rho$ if they observe $\underline{p}_L$.

c either. This can also easily be proven formally by applying Bayes' rule.
As no new information for a nonsearching consumer becomes available *if* they observe $p = \nu$, their beliefs update according to

$$\alpha_{t+1} = \rho\alpha_t + (1 - \rho)(1 - \alpha_t). \tag{4.9}$$

A consumer with initial belief ρ who observes ν in each subsequent period will thus have beliefs converging to 1/2 from above, that is $\alpha_{t+k} = 1/2$ for $k \to \infty$ according to Equation (4.9).
This yields another parameter restriction that is needed to derive equilibrium. It is assumed that $\hat{\alpha}$ will be such that, on the one hand, consumers with the most optimistic belief $1 - \rho$ will always search, and on the other hand, that consumers with the most pessimistic initial belief of ρ will never search if they observe a price of ν in every period that follows. This condition boils down to $\hat{\alpha} \in (1 - \rho, 1/2)$. Another interpretation of $\hat{\alpha} < 1/2$ is that critical consumers will never search if they believe that the high cost state is more likely than the low cost one.

Now how does asymmetric pricing emerge in the dynamic setting? First of all, a benchmark for prices can be defined. In a given period, the average price (and lowest price $\underline{p}$) will be lowest when the cost level is low ($c_t = c_L$) and all of the critical consumers search, i.e., $\mu_t^* = \overline{\mu} = 1 - \lambda_2$. Analogously, the average price (and lowest price $\underline{p}$) will be highest when the cost level is high ($c_t = c_H$) and none of the critical consumers search, i.e., $\mu_t^* = \underline{\mu} = \lambda_1$. Restricting the model parameters in a way that the search cost s_M of critical consumers is such that there must always be a positive fraction of critical consumers that does (not) search if $\alpha \in (0, 1)$, consumers will have full knowledge about the cost regime they live in if the search intensity is $\overline{\mu}$ or $\underline{\mu}$, as otherwise there would be heterogeneity among critical consumers. Because of this, it is possible to define full price adjustment under both cost states. Prices will have fully adjusted if it holds that $\mu_t^* = 1 - \lambda_2$ under c_L or $\mu_t^* = \lambda_1$ under c_H.

From this, it can be seen why there will be asymmetries in price adjustment. First, consider the case were high costs persist from period t onwards, i.e., having a cost realization of c_H for $t, t+1, t+2, ...$

In period t, the fraction $F_t(\hat{\alpha})$ of consumers that have beliefs $\alpha_t < \hat{\alpha}$ will search and observe the price $\underline{p}_H$. Thus they can immediately infer in which cost state they are (leading to an immediate partial price adjustment if costs were different before), and their initial beliefs for period $t+1$ will become $\rho > 1/2 > \hat{\alpha}$ according to the assumption made above that $\hat{\alpha} < 1/2$.

The fraction $1 - F_t(\hat{\alpha})$ of critical consumers that do not search in t will consist of consumers that happen to observe the low price $\underline{p}_H$ by chance (which is a fraction $\eta_t(\underline{p}_H)$ of nonsearching critical consumers) and consumers that only observe the high price ν. The nonsearching critical consumers that observe the low price can also immediately infer in which state they are (again leading to an immediate partial price adjustment in case of a cost regime change) and will adapt their beliefs for $t+1$ to $\rho > \hat{\alpha}$, just like the searching critical consumers. Solely the nonsearching critical consumers that observe ν cannot immediately infer in which state they are (averting full price adjustment in t if costs have changed) and also cannot correctly adapt their beliefs for $t+1$. Thus they will update their beliefs following 4.9, with $\alpha_{t+1} \in [1/2, \rho]$ as Equation (4.9) is approaching $1/2$ from above. Combining these observations, one can see that *every* critical consumer will have beliefs $\alpha_{t+1} > 1/2 > \hat{\alpha}$ in $t+1$, implying that

$$F_{t+1}(\hat{\alpha}) = 0 \quad \text{if } c_t = c_H. \tag{4.10}$$

Overall, a positive cost *shock* in period t will only have a partial price adjustment effect in period t (as some fraction of nonsearching critical consumers is unable to infer the true cost state in t), but will induce a full adjustment of prices in period $t+1$ because the search intensity μ will drop to $\underline{\mu}$ within one period as of Equation (4.10). However, this only holds if the cost remains high in period $t+1$, as only in this case beliefs and prices will not be affected by a chance observation of $\underline{p}_L$ by some nonsearching critical consumers in $t+1$.

In contrast, examine the case where a negative cost *shock* happens at period t, changing costs from c_H in period $t-1$ to c_L in periods $t, t+1, t+2, ...$ Like above, critical consumers in period t will be split into three subgroups. Those with a belief of $\alpha_t < \hat{\alpha}$ will search, observe $\underline{p}_L$, directly infer the true cost state from it and cause an immediate effect on prices. Their beliefs for $t+1$ will update to $\alpha_{t+1} = 1-\rho < 1/2$. The same happens for the fraction $\eta_t(\underline{p}_L)$ of critical consumers that do not search but observe the low price $\underline{p}_L$ by chance. The remaining $1-\eta_t(\underline{p}_L)$ nonsearching critical consumers that observe ν cannot infer new information about the current state (which results in no immediate effect on prices) and will adapt their beliefs for $t+1$ according to 4.9, with α_{t+1} remaining in $[1/2, \rho]$. As $1-\rho < \hat{\alpha}$, the former two subgroups will search in period $t+1$, while the latter subgroup will not. That is, $F_{t+1}(\hat{\alpha}) = F_t(\hat{\alpha}) + \eta_t(\underline{p}_L)[1 - F_t(\hat{\alpha})]$ or

$$F_{t+1}(\hat{\alpha}) = F_t(\hat{\alpha}) + \frac{\mu_t^* \beta}{k - \beta + \mu_t^* \beta}[1 - F_t(\hat{\alpha})]. \tag{4.11}$$

As it is assumed that there is an initial cost shock (costs are high in $t-1$), it follows from using 4.10[24] for $t = t-1$ (i.e., $F_t(\hat{\alpha}) = 0$ and $\mu_t^* = \underline{\mu}$) that

$$F_{t+1}(\hat{\alpha}) = \frac{\underline{\mu}\beta}{k - \beta + \underline{\mu}\beta} < 1 \quad \text{if } c_t = c_L, c_{t-1} = c_H. \tag{4.12}$$

[24]Recall that (4.10) doesn't depend on an initial cost shock to apply.

The above ratio is smaller than unity, implying that consumers' beliefs do not update as quickly (i.e., within one period) when there is a negative cost shock as when there is a positive one. Full price adjustment would only be reached if it held that $F_{t+1}(\hat{\alpha}) = 1$, which means that every critical consumer would search in $t+1$, i.e., it would have to hold that $\mu^*_{t+1} = \overline{\mu}$. Furthermore, even if the low cost state persists, there will always be some fraction of consumers that has not yet observed the low price $\underline{p}_L$. It can be shown via Equation (4.11) that a negative cost shock in period t which persists forever would need infinitely many periods to finally result in $F(\hat{\alpha}) = 1$, i.e., $\mu = \overline{\mu}$, resulting in full price adjustment.
It can also be seen that the magnitude of asymmetry in the first period after a cost shock will be quite large. As a full adjustment after a cost decrease implies that $F_{t+1}(\hat{\alpha}) = \frac{\underline{\mu}\beta}{k-\beta+\underline{\mu}\beta}$ should be equal to 1, the amount of adjustment will be small if $\underline{\mu}$ is small (or $k - \beta$ is large). The former is true when the fraction of consumers with very low search costs (who always search) is small, which seems to be a plausible assumption.

In summary, the model predicts a full adjustment of the search intensity (and hence, prices) to a positive cost shock within two periods, whereas a negative cost shock is predicted to induce a long adjustment period, with only a gradual increase in the search intensity and a gradual decrease in prices. The intuition behind this is that positive cost shocks are immediately learned by critical consumers who search, which directly influences their beliefs for the next period, making them nonsearchers. Critical consumers who do not already search either observe $\underline{p}_H$, infer from this that they are in the high cost state and stop searching as well, or observe ν and update their beliefs such that they stay nonsearchers in the next period. Thus none of the critical consumers will search in $t+1$. On the contrary, a negative cost shock can never have the effect that every critical consumer searches in the next period because those critical consumers who do not search in t *and* that do not happen to observe $\underline{p}_L$ in this period will keep nonsearching in $t+1$, as their beliefs only converge to $1/2 > \hat{\alpha}$ from above. In each subsequent period, some residual fraction of nonsearching critical consumers will never have observed $\underline{p}_L$ so far and will remain nonsearching in the period that follows. This is why there will only be a gradual increase in search intensity and slow decline in prices, compared to a positive cost shock.

Finally, it is worth noting that there is an infinite number of possible equilibrium paths, generated by different cost evolutions $\{c_t\}$. All of those would be characterized by different degrees of asymmetry, depending on how the cost states change. But it is not surprising that when examining the *expected* evolution path, adjustment asymmetry is obvious. This is because in expectation, each cost state (H or L) will last for some constant number of periods, say N.[25] The expected cost evolution will thus look like

$$\underbrace{LL...}_{N}\underbrace{HH...}_{N}\underbrace{LL...}_{N}... \,.$$

Following the argumentation from above, it is evident that the lowest price will be reached in the Nth period of state L, while the highest price will already be reached in the second period of state H. The larger the value of N is (implying a higher ρ), the more adjustment asymmetry will be present in the economy, on average.

In other words, the comparative statics for the persistence parameter ρ are such that a higher persistence of costs will lead to a more pronounced adjustment asym-

[25] The expected duration of one cost regime will be $\rho + \rho^2 + \rho^3 + \ldots = \frac{1}{1-\rho} > 2$.

metry of prices to cost shocks, at least on the expected evolution path of costs. More intuitively, this can be seen by defining the amount of asymmetry as the ratio of absolute price adjustments within one period after a negative and positive price shock (adjustment ratio). Then, it can be shown that this adjustment ratio will be smaller (implying more severe rockets and feathers) if ρ increases. The reason for this is that a larger ρ implies a longer expected duration of some cost state, resulting in a lower minimum price after an expected N periods of low costs. Then, the change in average prices after a positive cost shock will be larger, as positive cost shocks lead to a full adaption of prices after one period. On the other hand, the speed with which average prices fall after a negative cost shock does not depend on ρ. Overall, this means that the adjustment ratio will be smaller for large ρ, as the adjustment to positive cost shocks (which is in the denominator of the adjustment ratio) increases, while the adjustment to negative cost shocks (which is in the nominator of the adjustment ratio) stays the same.

A nice feature of this property is that it can be tested empirically. In markets where costs do not tend to fluctuate much, price adjustment should be more asymmetric than in markets with very instable input costs. Also, as this result was already obtained by the trigger sales model provided in Section 3, the property that asymmetric pricing should be more severe under a high persistence of costs gets further support.

The next parameter with interesting comparative statics is the fraction of shoppers λ_1. It can be shown that a smaller fraction of shoppers implies a longer adjustment process to negative cost shocks. I will again leave out a formal proof as the intuition behind this is quite striking. When a negative cost shock occurs in some period t and persists afterwards, a smaller λ_1 directly implies that a smaller fraction of consumers searches in each period, which leads to a slower adjustment of beliefs and prices in t and every period that follows. On the contrary, a positive cost shock will be learned by every critical consumer after one period, just as with a higher λ_1. Overall this leads to a higher degree of adjustment asymmetry.

Interestingly, this feature of the model is the exact opposite of the result obtained in the oligopolistic coordination model I provided in Section 3, where a higher fraction of shoppers in any submarket implied a higher likelihood of asymmetric pricing among firms. The fraction of shoppers λ can thus be considered as benchmark to test which model yields the better predictions: if a high fraction of informed consumers does in fact decrease the chance of asymmetric pricing in many markets, a trigger sales model in the form I provided would turn out to be an improper explanation for such a behavior.

Finally one can derive that an increase in the capacity of firms k compared to the measure of consumers β will also lead to a slower adjustment to negative cost shocks. The reason for this is that a smaller k/β will lead less firms to set the low price $\underline{p}$.[26] Hence in each period after a negative cost shock, fewer critical consumers (that do not search) will learn that they are in the low cost state by chance, resulting in a slower adjustment of beliefs and prices.

Discussion Like in any other economic model, the authors had to make simplifying assumptions. Probably the most unorthodox assumption of the model is that

[26] To see this, recall Equation (4.4) of the static model. As $\frac{\partial \eta(\underline{p})}{\partial k} = \frac{-\mu\beta}{(k-\beta+\mu\beta)^2} < 0$, fewer firms will set the low price $\underline{p}$ if k gets bigger compared to β.

there is not a discrete set of firms, but a continuum with measure one. The reason for this is mostly a technical one, i.e., it avoids having to deal with randomness and makes the derivation of equilibrium easier. I agree with the authors that the implications should be no different to a discrete model with, say, $n = 1000$ firms. However, as the number of firms reaches a level were collusion becomes possible (perhaps ten or less firms), I think the results will probably be inadequate and the model proposed in Section 3 might yield better predictions. But as the point of the model was not the description of behavior in an oligopolistic framework, this is just a minor point of consideration.

Another implication of the assumed continuum of firms is that searchers (who are able to observe every price quotation made by firms) effectively get to know infinitely many prices in each period. This seems to be quite an implausible property. To its defense, the authors show that there would be no explicit equilibrium solution for the case where each searching consumer only observes $n > 1$ prices (which would be a more realistic scenario).

Also I think that they are right in their view that the basic features of the model should carry over to a model with a finite number of observed price quotations in each period. The reason for this is that the key mechanism that drives price adjustment asymmetry in the model is that nonsearchers only observe one price quotation per period and can thus only slowly update their beliefs in case of a negative cost shock. In contrast, even if searching consumers only observe n instead of infinitely many price quotations, they will be able to update their beliefs much quicker than nonsearching consumers if a negative cost shock has happened. In fact, with a fraction $\eta(\underline{p})$ of firms charging the low price, the probability of observing at least one low price $\underline{p}$ when observing n random prices is $1 - [1 - \eta(\underline{p})]^n$. This value quickly converges to one for increasing n.

The last unconventional assumption is that each firm has a capacity constraint of exactly k units. While it is probably true that there are some (possibly pretty loose) constraints in most retail markets, the model implies that these constraints play a crucial role for reaching equilibrium, which is, in my opinion, a bit far-fetched.[27] The authors address part of this defect by relaxing the model assumptions to allow for two different capacity constraints k_1 and k_2. They can show that a two-point distribution of prices will also emerge in this case and that asymmetric pricing is robust to heterogeneous capacity constraints (at least for two different constraints).

Overall, none of the assumptions of the model seem to be indefensible, although the property of binding capacity constraints for each firm that sets the low price $\underline{p}$ seems to be quite unrealistic. Also it has to be emphasized that Yang & Ye's model only addresses the case of markets with a large number of competing firms. Once collusion gets more probable and the number of firms decreases, the model's implications might become inaccurate.

As a final remark, it has to be said that the authors are among the first who provide a theoretical explanation of why asymmetric price adjustment might naturally emerge in competitive markets. If some of their predictions can be verified empirically, their contribution should have a large impact on contemporary economic theory, especially the theory on consumer search and optimal pricing under competition.

[27] E.g. a big gasoline retailer will usually not have a binding capacity constraint.

4.2 A Model of Consumer Search Without Learning

Model Setup and Equilibrium of the Static Game In this subsection, I will present a model suggested by Tappata (2009) that explains asymmetric price adjustment but doesn't rely on a learning process to derive this asymmetry. In some way, Tappata's model can thus be considered as a more elemental version of Yang & Ye's work. The price for this is that a gradual adaption of prices cannot be modeled, meaning that prices always fully adjust within two periods.

Tappata uses a model based on nonsequential consumer search in a competitive market with rational agents to explain asymmetric price adjustment by firms to random cost shocks. In order to do so, the author points out that the demand elasticities of consumers must be a function of previous cost realizations.[28] If this was not the case, firms would have no incentive to differentiate between positive and negative cost shocks when determining their optimal output prices. Even in the case of imperfectly informed consumers with positive search costs, the resulting equilibrium *price dispersion* (see Varian, 1980) will result in a new optimal markup that doesn't discriminate between positive and negative cost shocks.

So what is the key mechanism that drives asymmetry in Tappata's model? Like in many prominent search models, firms choose some optimal price distribution $G^*(.|\mu, c)$ that is dependent on the search intensity μ of consumers and the firms' cost level c. Similar to Varian's model, firms will choose less dispersed prices under high than under low costs because their pricing range, i.e., the difference between marginal costs and the monopoly price, gets smaller under high costs. As

[28] Or, as in Yang and Ye (2008), beliefs about previous cost realizations.

consumers are rational, they adapt to this behavior by searching less when they expect prices to be high.
Here comes another feature into play that was already used in Yang & Ye's model, namely that cost shocks can be time dependent. Again, this means that there can be some persistence parameter $\rho > 1/2$ that defines the likelihood of some cost state carrying over to the next period. So if $\rho > 1/2$ (and consumers know this), search will be low if the previous period's costs were high and high when the previous period's costs were low. This is exactly what constitutes demand elasticities that are dependent on *previous* cost realizations, but not current ones.
An unexpected change from low to high costs will thus lead firms to quickly increase their prices because consumers will search a lot. In contrast, a change from high to low costs will have a smaller effect because less consumers will search. Overall this implies asymmetric price adjustment to costs.

Having outlined the intuition behind Tappata's model, it is time to briefly describe the model setup in the static framework. The industry constitutes of $n = 2$ firms that sell a homogeneous good and compete through prices. At the beginning, *both* firms are randomly assigned either low ($c = c_L$) or high ($c = c_H$) production costs, with the probability of both firms having costs of c_H being α and both firms having costs of c_L being $1 - \alpha$.

While the model can be fully extended to $n > 2$ firms in the market, with most of its predictions being robust to the introduction of a multitude of firms, I will not discuss the model for the case where $n > 2$. However, I will briefly point out some of its implications for increased n in the concluding discussion.
But no matter what n is, the first crucial difference to Yang & Ye's model becomes apparent. Instead of a continuum of firms (that creates issues like infinitely many

price observations per searcher), there is only a finite (and possibly small) number of firms in the market.

Like in most search models, there is a continuum of consumers (this time with measure one) having a unit demand and a reservation price ν. Before consumers buy, they decide whether to search or not by the protocol of nonsequential search. If they search, they get to know all (i.e., for $n = 2$, both) prices set by the firms and shop at the firm with the lowest price. If they do not search, they shop at a random firm.
Again, some fraction λ of consumers called *shoppers* has zero search cost, while all other (i.e., $1 - \lambda$) consumers are called *nonshoppers* and have positive search costs drawn from a continuous distribution, with $s_i \in S = [0, \overline{s}]$. $H(s)$ denotes the fraction of nonshoppers that have search costs smaller than s.[29]

Depending on the expectation of consumers about the firms' cost level and the resulting price dispersion, only those nonshoppers will search that have search costs lower than the expected gains from search, resulting in some search intensity μ. At the same time, firms anticipate this search intensity and maximize their profits by setting prices accordingly. Equilibrium can thus again be characterized by some equilibrium search intensity μ^* and an equilibrium distribution of prices $G^*(.|\mu^*, c)$.

From the above assumptions, it is easy to calculate the profit of a firm charging p_j, facing a search intensity of μ, costs of c and the price of the other firm p_{-j}. It holds that

$$\pi_j(p_j, p_{-j}; c; \mu) = (p_j - c)[\frac{1+\mu}{2} I_{\{p_j < p_{-j}\}} + \frac{1}{2} I_{\{p_j = p_{-j}\}} + \frac{1-\mu}{2} I_{\{p_j > p_{-j}\}}] \quad (4.13)$$

[29]This distribution of search costs is assumed to be public knowledge.

with $I_{\{X\}} = 1$ if X is true, and 0 if not. For example, if $p_j < p_{-j}$, it follows that $\pi_j(p_j, p_{-j}; c; \mu) = (p_j - c)(\frac{1+\mu}{2}) = (p_j - c)(\mu + \frac{1-\mu}{2})$, i.e., the full measure of shoppers and half of the measure of nonshoppers is served for the price of p_j.

Then, the utility of some consumer i can be defined as the difference of their reservation price ν to the expected price they have to pay (including search costs in case they search).

To solve for a (symmetric) Bayesian Nash equilibrium where the firms' pricing strategy $G(.|c)$ is a best response to the consumers' search decision μ and the consumers' search decision is a best response to the firm's pricing decision, it is useful to start with the profit function of firms. Using an undercutting argument similar to the one provided in Subsection 3.2 and Varian (1980), one can show that that there can be no symmetric equilibrium in pure strategies if $0 < \mu < 1$.

Solving for a mixed strategy equilibrium where both firms price according to some probability distribution $F(p)$, one can show that

$$F(p) = 1 - \frac{1-\mu}{2\mu}\left(\frac{\nu - p}{p - c}\right) \tag{4.14}$$

for any p in a range of $[\underline{p}, \nu]$, with

$$\underline{p} = c + \left(\frac{1-\mu}{1+\mu}\right)(\nu - c). \tag{4.15}$$

For a proof that (4.14) and (4.15) do in fact constitute a unique Nash Best Response to μ, see Varian (1980).

It can easily be seen that $\underline{p}$ is negatively related to the search intensity μ, i.e., firms' pricing range increases when consumers search more. The reason for this is that more informed consumers imply smaller profits for firms, extending the range of prices that firms can charge while remaining indifferent between them.

Also, it can be seen that $F(p, \mu'; c)$ is bigger than $F(p, \mu; c)$ for every $\mu' > \mu$. That is, the probability of observing some price smaller than p is always larger for higher search intensities $\mu' > \mu$. This can be explained by the fact that higher search intensities make it more profitable for firms to try to steal informed consumers from their competition by charging the lowest price in the market.

Next, the search decision of individuals (who are not capable of influencing μ because there are infinitely many consumers) will be examined. Given firms' best response to some search intensity μ that was derived above, a consumer's expected benefit from search is their expected price reduction given μ, i.e., $E[p - p_{min}|\mu]$. Tappata shows that

$$E[p - p_{min}|\mu] = (\nu - E[c])\frac{1-\mu}{2\mu^2}\left[\log\left(\frac{1+\mu}{1-\mu}\right) - 2\mu\right] \tag{4.16}$$

with $E[c] = \alpha c_H + (1-\alpha)c_L$.

The above expression can also be interpreted as price dispersion. The higher this dispersion gets, the more consumers will decide to search (given the distribution of search costs $H(s)$) because more and more consumers will have an expected benefit from search that exceeds their search costs. According to Tappata, Equation (4.16) can be shown to have an interior maximum at some point $\hat{\mu}$ with $0 < \hat{\mu} < 1$, meaning that beyond some optimal search intensity, the expected benefit of search begins to diminish.

Like in other search models, a consumer will decide to search when the benefits of search exceed their search cost. It is clear that the fraction λ of consumers with zero search cost (shoppers) will always search. In contrast, consumers with search cost greater than $\nu - p^*$ will never search, as their search cost will always exceed their benefits from search. That is, a fraction $(1-\lambda)(1-H(\nu-p^*))$ will strictly prefer to shop at a random firm in equilibrium.

The search costs of an indifferent consumer $\tilde{s}$ must be equal to their expected benefits from search, i.e., their net benefit from search must be zero. That is,

$$\tilde{s} = (\nu - E[c])\frac{1-\mu}{2\mu^2}\left[\log\left(\frac{1+\mu}{1-\mu}\right) - 2\mu\right] \tag{4.17}$$

Having defined $\tilde{s}$, every consumer with search costs less than $\tilde{s}$ will search, while every consumer with higher search costs will not. To be precise, it must hold that

$$\mu = \lambda + (1-\lambda)H(\tilde{s}). \tag{4.18}$$

By combining equations (4.17) and (4.18), an equilibrium solution for $\tilde{s}$ and μ can be found. This market equilibrium depends on the number of shoppers λ, the distribution of search costs $H(s)$ and the parameters that influence the gains from search. One problem is that it doesn't have to be unique, meaning that there can be multiple $\tilde{s}$ and μ where a marginal consumer is indifferent between searching and not. However, it can be shown that the equilibrium must be unique if $\lambda > \hat{\mu}$.[30]

One example for a unique equilibrium can be seen in Figure 4.1. Here, search costs ($s_i \in [0, \overline{s}]$) are assumed to be uniformly distributed among nonshoppers, implying that some marginal nonshopping consumer that has not yet searched will have search costs of $(\mu - \lambda)\frac{\overline{s}}{1-\lambda}$ when μ consumers have already decided to search. The equilibrium levels of μ and $\tilde{s}$ can easily be determined by the intersection of the gains from search and search cost lines. The blue curve depicts the expected gains from search under higher expected costs, i.e., higher α.

[30] For another parameter condition that implies unique equilibria, see Tappata (2009, p. 679).

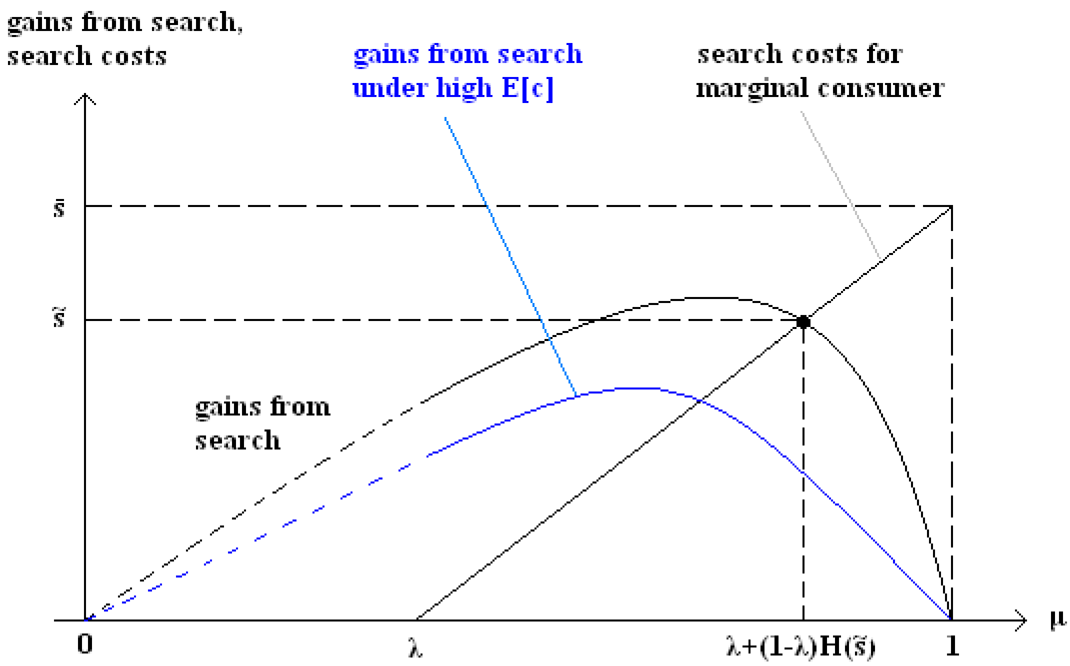

Figure 4.1 Example for a unique equilibrium in Tappata's model. Source: Modification of Tappata (2009, p. 679)

As we see, equilibrium constitutes of some optimal search intensity μ^* that is a best response to firms' expected price dispersion and an optimal price dispersion $E[p - p_{min}|\mu^*]$ that is a best response to this optimal search intensity μ^*. The higher the equilibrium search intensity, the more dispersed prices are.

Recalling Equation (4.15), it is important to recognize that the range of prices firms will choose from is smaller with high production costs c. As a result, the price dispersion (and thus, benefits from search) are smaller under high production costs, implying that consumers who *expect* costs to be high will search less.[31]
Because the expectation of consumers can be wrong and firms have an informational advantage over them by getting to know the true production costs without delay, firms' optimal response to cost shocks will be different for positive and negative shocks. As mentioned earlier, a sudden change from low to high production costs in period t will result in a 'too high' search intensity of consumers in period $t + 1$ (compared to if they knew in which cost state they are), leading firms to increase prices more than it would be optimal under perfect information. In contrast, a sudden dump in production costs has the effect that 'too few' consumers will search in the next period, implying that firms' optimal strategy is to reduce prices less than it would be optimal in a world with perfect information. This is exactly the mechanism that drives asymmetric pricing in Tappata's model. In what follows, I will derive this asymmetry for the dynamic game.

[31] This can be seen by considering Equation (4.16) for increasing α. As the gains from search get smaller for each search intensity μ if $E[c]$ increases, the search costs of the indifferent consumer $\tilde{s}$ are smaller for high α, implying less search intensity μ under pessimistic consumers. This can also be seen graphically by considering the blue curve in Figure 4.1.

Equilibrium of the Dynamic Game Here, I will proceed to extend the static version of Tappata's model to a dynamic framework and show that asymmetric price adjustment does emerge if production costs are persistent, i.e., cost shocks are *not* independent and identically distributed. In order to do so, the static game is repeated over time, with $t = 1, 2, ...$ At the beginning of each period, nature assigns high production costs with probability α and low production costs with probability $1-\alpha$ to both firms. Then, consumers get to know the previous period's cost realization and can update their beliefs accordingly. Finally firms maximize their utility by choosing some price distribution that maximizes their profits, while anticipating the search decision of consumers.
It is clear that in this setup, prices can change because of two reasons. First of all, firms will adapt their pricing strategies to cost shocks. But the more important feature of the model is that the distribution of prices will also change when consumers' priors relative to the true state of the economy change. This means that the demand elasticities of consumers will depend on previous cost realizations, allowing firms to react asymmetrically to cost shocks.

Similarly to the model I provided in Section 3 and the model discussed by Yang and Ye, the probability α of the high cost state will be ρ if there were high costs in the last period and $1-\rho$ if there were low costs in the last period, with $0 < \rho < 1$. As consumers always get to know the previous period's cost realizations and adapt their beliefs immediately, there are only four different states of the economy to consider. If one again denotes the low cost state by L and high cost state by H, these states (c_{t-1}, c_t) are $\{LL, LH, HL, HH\}$. Also, the search intensity of consumers will only depend on the previous period's cost state.

In this setting, it is easy to verify that it will take only two periods for prices to fully adjust to some cost shock. In the period where the cost shock happens (say t), consumers will have some (distorted) prior about the current cost realization and firms will adapt to the cost change by modifying their price distribution accordingly, both taking into account the new cost level and consumers' beliefs. But if the new cost state persists, consumers get to know the true state in $t+1$, adjust their search intensity and lead firms to price optimally given the (now correct) beliefs of consumers. Thus, in order for pricing asymmetry to exist, asymmetry must be found in the very period where the cost shock happens (i.e., in the first period with the new cost state).

More explicitly, adjustment asymmetry will be prevalent if the expected absolute change of the average price in the market is bigger for positive cost shocks than for negative ones. For this, it has to be considered that a positive cost shock can only happen if the previous period's costs were low, i.e., the cost state must have been LL or HL and change to LH. Analogously, a negative cost shock can only happen if the economy's state was LH or HH and changes to HL. Overall, Tappata is able to determine the difference of the expected absolute change of the average market price after a positive cost shock compared to a negative cost shock as

$$E[|\Delta p|\ |\Delta c > 0] - E[|\Delta p|\ |\Delta c < 0] = \frac{-\rho(1-\rho)}{2}[(p_{HH} - p_{HL}) - (p_{LH} - p_{LL})] \quad (4.19)$$

where p_{xy} denotes the average market price under the respective cost state xy.

The usual case of a slower adaption to negative cost shocks is thus found if the above expression is larger than zero. It is possible to decompose the sign of (4.19) such that it will only depend on the sign of the following: $\frac{-\rho(1-\rho)}{2}(c_H - c_L)$ times

> i) [t]he effect of previous cost realization on consumers' priors, ii) the effect of those priors on the equilibrium search intensity, and iii) the effect of the search intensity on the cost pass-through. (Tappata, 2009, p. 682)

As mentioned earlier, the effect of consumers' priors α on the search intensity μ is clearly negative ($\frac{\partial \mu}{\partial \alpha} < 0$) while it is obvious that the effect of the change of previous period's costs (compared to the period before that) on consumers priors is positive ($\frac{\Delta \alpha}{\Delta c_{t-1}} > 0$) if costs are persistent, i.e., $\rho > 1/2$. What remains is the effect of the search intensity on the cost pass-through. Using expected market prices for given cost realizations c and consumers' priors α, Tappata is able to derive this equilibrium pass-through of costs. To be precise, he shows that

$$\frac{\partial E(p|c)}{\partial c} = 1 - \frac{1-\mu}{2\mu} \log \left[\frac{1+\mu}{1-\mu} \right] \geq 0. \tag{4.20}$$

As the pass-through of costs (given some cost level c) is dependent on the expected elasticity of demand faced by firms and this elasticity of demand increases with the search intensity μ, the cost pass-through will positively depend on μ (this can easily be shown by differentiating (4.20) with respect to μ). Using L'Hôpital's rule, one can also show that the pass-through for perfect competition ($\mu = 1$) is one and the pass-through for the monopoly case ($\mu = 0$) is zero, as expected.

Because of the above results, the sign of Equation (4.19) will solely depend on $\frac{-\rho(1-\rho)}{2}$ (that is, the process that drives α). Because the sign of this expression is negative, the overall sign of (4.19) will be positive (negative, positive, positive and negative yields positive).

In other words, adjustment asymmetry to cost shocks in the traditional sense can be found in the dynamic setting. Again, the intuition for this is that firms face higher demand elasticities in periods where a positive cost shock has happened than in periods where a negative cost shock has. This is because high costs in the previous period (i.e., a negative price shock) lead consumers to search too little as they expect costs to remain high. In contrast, consumers will search too much when a positive cost shock has happened, leading firms to substantially increase prices. Note that while there is pattern asymmetry in this model (with a different immediate price reaction to cost shocks), no timing or amount asymmetry can be found. After two periods, prices and beliefs will always return to the long-run equilibrium values.

Summary and Discussion Tappata's model is able to explain an asymmetric response of retail prices to random cost shocks if cost shocks are persistent. In order to do so, Tappata models the demand elasticities faced by firms such that they depend on past cost realizations observed by consumers. As consumers always observe the true cost realizations with a delay of one period, Tappata's model is in a way more elementary than Yang & Ye's. While in citetyang, it would (in principle) be necessary to trace the past beliefs of a continuum of consumers for all elapsed periods,[32] Tappata only needs to distinguish between four states of the economy. In each state, the beliefs of consumers are uniquely determined and will either be correct (if no cost shock took place at the beginning of the period) or distorted. This leads firms to adapt differently to positive and negative cost shocks as the search intensity of consumers and thus the demand elasticities faced by firms will depend on consumers' beliefs about present costs, which will be distorted if a cost shock happened.

I want to emphasize one more time that this simplification in Tappata's model comes at the price that cost shocks are always absorbed within two periods. While Yang & Ye are able to model a slow adaption of output prices to a negative cost (while positive shocks are fully priced in after one period), prices will immediately

[32] Yang & Ye overcome this technical problem by assuming a continuum of firms, which in turn implies some theoretical objections (outlined in the discussion of the last subsection).

jump to their stable values after a one time shock in Tappata's economy in the following period. Empirically, the latter adjustment pattern cannot be confirmed.

There are some other differences between Yang & Ye's and Tappata's model. First, while Yang & Ye use a continuum of firms to derive equilibrium in their paper, Tappata focuses on the case of only two firms. Although his results carry over to the case with more than two firms in the market, some of the implications of this are quite awkward. While an increased number of firms is usually associated with a better functioning of markets (reaching perfect competition in the limit), it can be shown that an increased number of firms in the model economy actually raises average prices for given search intensities μ.
The expected cost pass-through of prices can also be proven to diminish for increasing n, implying that atomistic markets will have a very slow adaption to cost shocks for values of μ that are not close to unity. Depending on how the exact distribution of search costs for consumers looks like, this implies that the model makes a prediction for competitive markets that cannot be confirmed by reality if there is even a small fraction of consumers that finds it unprofitable to search in any case.

Another discrepancy between the models is that Tappata uses no capacity constraints of firms to derive his results. An important objection to Yang & Ye's model was that each firm pricing the (for each state uniquely determined) low price $\underline{p}$ had to be capacity constrained. Clearly this is quite an implausible assumption for most retail markets. Another problem was that each firm was assumed to have the same capacity constraint, which is an unrealistic assumption. Tappata in turn needs no capacity constraints at all, which seems clearly superior.

Possible extensions to Tappata's model would be the consideration of oligopolistic behavior among firms for small n, an imperfect learning process for consumers or the incorporation of dynamic learning that is influenced by past search decisions (consumers who have searched in the past might recognize cost levels more accurately than consumers who have not). Also, a mechanism could be conceived that ensures that average prices will not rise once the number of firms rises beyond some point.

Overall, Tappata's model can be understood as a concise theoretical work that is able to explain asymmetric pricing in competitive markets while only relying on few limiting assumptions. Although some of the model's predictions are not in accordance to the behavior of real markets, Tappata provides a very interesting starting point for future research in the area of asymmetric pricing caused by consumer search. In contrast, Yang and Ye rely on more limiting assumptions, but are able to achieve results that closer match current empirical evidence.

4.3 Other Consumer Search Models

4.3.1 A Reference Price Search Model with Asymmetric Pricing

In a yet unpublished working paper, Lewis (2009) suggests a *reference price* consumer search model that results in asymmetric pricing. The basic mechanism which is employed in Lewis' model is as follows. Consumers observe one random price in the market for free, but have to pay positive search cost if they want to see another price quotation.
Most importantly, unlike in the consumers search models that were discussed above, consumers are not acting fully rationally: they do not know the random process with which firms' costs are determined and thus base their expectations of current prices on the prices they have observed in previous periods, using some prior distribution of prices that can be different from their actual distribution. In the model, expectations need to be biased in order to produce asymmetric pricing. Given these expectations, both consumers and firms act rationally. However, wrong expectations lead consumers to either search too much or search too little when costs change.

In particular, consumers who observe a price that is lower than the price they paid in the last period will assume that there is only a small chance of finding an even lower price when searching, implying that fewer consumers will search. Thus, firms can get away with higher margins because there will be less competition among them (the fewer consumers search, the more incentive a firm has to price high, as it doesn't care about the small fraction of searching customers it loses to a firm that prices lower). Basically, if costs drop significantly, firms only need to reduce prices a little to keep consumers from searching, and thus negative costs shocks are transmitted slowly to output prices.

On the other hand, if there is a significant positive cost shock, firms will have to adapt their prices instantaneously because margins are squeezed and would become negative otherwise. As no consumer would expect such a big rise in costs, every consumer would choose to search after observing the first price, implying that the only price firms can charge in equilibrium are the marginal costs firms are facing, meaning that a full information Bertrand-type scenario results.

Overall, asymmetric pricing has to emerge: output prices react immediately to 'large' (in relation to margins) positive cost shocks, but slowly adapt to large negative cost shocks. If cost shocks are small relative to margins, there will be a slow decline of market prices no matter whether costs rise or fall.

Lewis' model of asymmetric pricing has one major implication. As mentioned above, his model predicts that prices should only follow changes in marginal costs when costs have risen significantly (such that firms need to price higher in order to avoid making losses) or have fallen slightly. If costs drop excessively, firms will choose to decrease them just as much to ensure that no consumer searches, generating high margins. If costs rise slightly but margins are very high, firms will opt to further decrease their selling prices in order to avoid generating search.

In other words, prices will only be sensitive to cost shocks if margins are small. If margins are small and a positive cost shock happens, firms need to adapt their output prices immediately in order to retain a positive price-cost margin. If margins are small and costs decline, prices will fall slightly (in concordance with costs) because firms try to prevent search.
However, if margins are high and a (large) negative *or* (small) positive cost shock happens, firms will optimally react to consumers biased expectations by slightly reducing their prices in an attempt to discourage search.

Lewis' main finding hence is that margin size, rather than the sign of a cost shock, is the main driving force behind the speed of price transmission to cost changes. The phenomenon of asymmetric pricing is only apparent because small margins (implying a quick response to cost shocks) usually correlate with cost increases, while high margins (implying slow price adjustment) correlate with cost decreases.

Interestingly, Lewis' model predicts similar dynamics to the dynamics obtained by the collusive trigger sales model proposed by Borenstein et al. (1997), which I tried to formalize in Section 3. If costs decrease significantly in the model (i.e., costs drop from c_H to c_L), firms start colluding on the old equilibrium price, with average prices decreasing slowly because of random demand shocks. As a consequence, a significant negative cost shock leads to high margins and a slow transmission of prices, just as in the model of Lewis.
On the other hand, if costs rise by a significant amount, margins drop to zero and firms must behave competitively, pricing at cost without delay. This means that a large increase in costs has an immediate effect on prices and goes hand in hand with low margins, again just like in Lewis' model.

It has to be noted, however, that my model is only able to differentiate between two cost states, which was necessary to allow for a unique focal point on which firms can collude once costs drop. Therefore, no statement can be made about what the effect of a small cost decrease under high margins would be, as in the model, high margins can only be found if costs are already low. Analogously, low margins coupled with a significant cost increase are impossible too, as zero margins can only occur if costs are already high. However, low margins coupled with a cost increase should naturally lead to a rise in prices, as margins would become negative otherwise.

Also, Lewis finds that the consumer search models of Yang and Ye (2008) and Tappata (2009) contradict some of the implications of his model, at least if they are generalized to an arbitrary amount of different cost states:

> In the Tappata (2008) and Yang and Ye (2008) models, more search and faster price response occur when dispersion is largest, and dispersion is inversely related to marginal cost. Therefore, a more generalized version of these models would predict that prices respond more quickly to cost changes when margins are high than when they are low. This contradicts Prediction 2 of the reference price search model. (Lewis, 2009, p.17)

Now that I've compared the implications of Lewis' model to those of the models explained above, I want to briefly address two of its problematic issues. First, as mentioned before, the consumers in his model are not perfectly rational in the sense that consumers' expectations always need to be biased in order to generate equilibria where asymmetric pricing occurs. Such an assumption is not needed by the consumer search models discussed before.

Second, Lewis assumes that firms are myopic in the sense that in each period, they try to maximize their profit without taking into account supergame strategies. This assumption effectively destroys any possibility of collusion in the model. But models where consumers have wrong expectations would certainly be prone for collusive strategies, which is ignored completely.

Overall, while Lewis' model provides a plausible alternative mechanism for asymmetric pricing, some of its assumptions are probably not justified for many markets. On the other hand, it generates empirical implications that can easily be tested. This is also done in Lewis (2009), where the author is able to confirm most of his model's predictions. While this doesn't mean that other models of asymmetric pricing are wrong or inferior to the one of Lewis, his empirical research points out several key implications any model of asymmetric pricing should be able to provide. Most importantly, Lewis finds that the speed of price adjustment to costs is typically dependent on margin size. This in turn should help researchers in the development of better or more general models of asymmetric pricing.

4.3.2 Asymmetric Pricing caused by Sticky Consumer Prices

Cabral and Fishman (2008) provide a model of sticky consumer prices when input prices are sticky. Under certain conditions, they show that the model also implies asymmetric price transmission in the traditional sense (i.e., a faster response to positive cost shocks), however only if input cost changes are small.

The underlying principle of the model is as follows: if a firm's input cost increases by a small amount, it might be optimal for the firm to refrain from increasing its output price to the new optimal level, as an increased price might lead consumers to search. If the expected profit loss caused by searching consumers (buying at another firm) exceeds the expected additional profit caused by adapting prices to the new optimal level, firms should not increase prices.
In contrast, if costs decrease by a small amount, firms have no incentive to cut prices because consumers optimally decide not to search after they observe unchanged prices. This is the basic mechanism why prices are sticky in the model.

Asymmetric price transmission finally occurs if cost changes are correlated across competing firms. In principle, if costs are sufficiently correlated, the signal of a slightly increased price tells consumers that market prices have likely increased in general (due to a positive cost shock), implying that consumers might find it optimal not to search. Because of this, firms can immediately react to positive cost shocks by increasing their selling prices. On the other hand, the dynamics of the model ensure that a small decrease in costs bears no incentive for firms to reduce their prices, as consumers will refrain from search anyway if prices don't change at all.

It can be shown though that no large price increase can be found in the model, which is an implication that is very different from those of all other models discussed in this thesis. Also, price decreases are typically large in the model, which is again different to the results of the other models discussed.

Overall, the model of Cabral and Fishman predicts the following behavior, which can also be tested empirically:

- Small cost decreases are transmitted with lag whereas small cost increases are transmitted to prices instantaneously. Large cost decreases or large cost increases do not yield asymmetric price adjustment.
- Cost changes and price changes have a higher correlation if costs are increasing than if costs are decreasing.
- Price decreases are typically less frequent than price increases, but price decreases tend to be larger in magnitude.
- Prices decreases are only less frequent than price increases if costs change by a small amount. If costs change significantly, no such asymmetry is found.
- Asymmetric price adjustment is only found if costs are (very) sticky.

Out of those, only implications two and five are shared by most of the models that have been discussed. While Cabral and Fishman provide empirical support for some of the other results, I think that implications one and four are somewhat awkward. It seems unlikely that adjustment asymmetry should only be found if cost changes are small. Intuitively, the other way round seems more plausible: a large cost increase needs to be transmitted immediately to prices in order for firms to avoid making losses whereas a large cost decrease can be exploited to achieve high margins for several periods.

Implication three is quite interesting though because it seems to resemble the pricing behavior in some markets where prices fluctuate much. For example, retail gasoline stations are often said to behave as follows: in the morning, prices are lowest but they continuously increase by small amounts throughout the day. Significant price drops are typically only found at early morning or after weekends.

While I have no empirical support for this claim, if retail gasoline stations and other firms do in fact act like this, the model of Cabral and Fishman does a good job to capture the effect. Also, all other models I have discussed predict the opposite behavior: price increases are typically large and non-recurring, whereas prices decreases are small and recurring until minimum costs are reached eventually, given some persistent lower cost level.

Thus, the model of Cabral and Fishman, although not specifically designed to yield asymmetric price transmission, provides a different mechanism to explain asymmetric pricing which is worth considering for future research.

5 Miscellaneous Sources for Asymmetric Price Adjustment

Now that several models of the two most common explanations for asymmetric pricing, namely oligopolistic coordination and imperfect consumer search, have been discussed, I want to briefly present some alternative ideas to model the phenomenon. In this last theoretical section of my thesis, I will describe the intuition behind three different and unrelated mechanism that might lead to asymmetric pricing.

In Subsection 5.1, I will start by presenting an idea of Borenstein et al. (1997), where asymmetric pricing is suggested to be the result of a slow adaption of production to cost shocks. In Subsection 5.2, I will point out a model of Ball and Mankiw (1994), in which asymmetric price adjustment is caused by asymmetric menu costs under inflation. Finally and probably most interestingly, I will discuss a model of Eckert (2002) in Subsection 5.3. There, pricing patterns that resemble rockets and feathers emerge in highly competitive markets as the result of price wars in which firms battle over market share. Also, as will be seen, an extension of the model considered could lead to 'true' asymmetric pricing.

5.1 Asymmetric Pricing due to Lags in Adjustment of Production and Finite Inventories

One alternative mechanism that could result in asymmetric was already given by Borenstein et al. (1997). In their second hypothesis, they argue that

> [p]roduction lags and finite inventories of gasoline imply that negative shocks to the future optimal gasoline consumption path can be accomodated more quickly than positive shocks. (Borenstein et al., 1997, p. 327)

While their paper concentrates on the gasoline market, this concept could also be extended to other branches.

In short, if a severe supply shock happens, asymmetric price adjustment could result because it might take some time in order to adapt production to future optimal consumption. To see this, consider a significant positive supply shock first, e.g. because production costs drop sharply. Then, it is clear that future output prices will have to sink eventually. However, prices cannot fall as much in the short run because inventories are finite and it will take a couple of days or weeks to alter production to increase output. This implies that negative cost shocks are transmitted slowly to selling prices.

On the other hand, if a severe negative supply shock happens (e.g., because production costs increase sharply), firms cannot profitably decrease current production. However, they can simply choose to raise prices immediately (building up inventories) in an attempt to cover their increased expenses later. This leads to a quick rise in prices once production costs rise.

Overall, an asymmetric price adjustment to (significant) cost changes will be found.

The above mechanism is certainly interesting because it differs greatly from all other explications of rockets and feathers presented in this thesis. I think that for markets with severe production bottlenecks, high market concentration and sig-

nificant supply or demand shocks, it could turn out to be one of the main driving forces behind asymmetric price adjustment.
However, most retail markets are characterized by different conditions: in many cases, there are no production bottlenecks and retailers can choose to buy virtually unlimited amounts of the commodity they obtain from wholesale. Also, severe supply or demand shocks are likely uncommon in many markets. As a result, the hypothesis seems worth considering for future research, but is unlikely to describe asymmetric pricing in every or a majority of markets.

5.2 Asymmetric Pricing due to Asymmetric Menu Costs

Ball and Mankiw (1994) propose a partial equilibrium model of asymmetric price adjustment that is caused by asymmetric menu costs under positive trend inflation. The intuition behind their model is as follows. If inflation is positive and a firm wishes to decrease its price relative to the aggregate price level (e.g. because of a negative cost shock), it can simply wait until its desired relative price level is reached. This is possible because positive trend inflation increases the aggregate price level and thus reduces relative prices over time. If the firm wishes to increase its relative price level, however, it has to offset the inflation in the market *and* the gap to its new desired relative price. It cannot wait until its desired relative price level is reached because, in the absence of another shock to its desired price, this gap would become wider and wider.

If there are no menu costs in the model, firms would simply set their relative price equal to their desired price in every period. But if there are menu costs which are large enough to have an effect on the choice of whether to change prices or not, asymmetric pricing will occur: price adjustment will be profitable more frequently when there is a positive shock to the desired relative price compared to a negative shock, and if prices are adjusted, adjustment will be larger under positive shocks.

While the simplicity of this model is certainly appealing, I don't believe that it is able to capture the true mechanism that leads to asymmetric pricing. The main reason for this is that inflation rates are usually very small compared to changes in input and output prices. Costs and other determinants of desired relative prices tend to fluctuate so much, in relation to their levels, that a typical single-digit inflation rate could never hope to have a significant effect on the adjustment of output prices.
Also, menu costs are likely very small in the great majority of retail markets. For example, a gasoline retail station that wants to adjust its selling price is probably not worried about the cost this change of price creates. In fact, I think that menu costs (in the traditional sense) are much too small to have an influence on the decision of changing prices in most markets were input prices fluctuate regularly. Thus, in sum, I don't think that the model of Ball and Mankiw (1994) will play a crucial role in identifying the sources of asymmetric pricing.

5.3 Edgeworth Price Cycles Resembling Asymmetric Pricing

Finally, a very interesting alternative to rockets and feathers was provided by Eckert (2002). In his model, a pattern similar to asymmetric pricing emerges although prices move more or less independently from costs. In the model, firms' price setting is characterized by two different regimes,

> an undercutting regime, in which firms battle over market share, and an increasing regime, in which firms cease battling for the market and instead choose to restore temporarily high prices. (Eckert, 2002, p. 64)

In other words, firms slightly undercut each other in the former regime, leading to a slow decline of average market prices until prices are close to marginal cost. Once they are sufficiently close to it, firms stop undercutting, price at cost and begin a 'war of attrition'. During this war, firms randomize between continuing to price at cost or increasing prices to a much higher level. Once one of the firms does so, it loses most or all of its market share for the next period, but the other firms in the market quickly follow to price at the high level. This is how the increasing regime is characterized. Once all prices have reached the highest level, a new undercutting regime starts.

Overall, the pattern of price movements that emerges (see Figure 5.1) can get confused with asymmetric pricing if input costs are unobservable: price decreases are small and the lowest price in the cycle is approached slowly whereas price increases are large, almost simultaneous and prices jump to the highest price in the cycle.

As argued above, Edgeworth price cycles can look like rockets and feathers even though prices move independently from costs.[33] Also, following the author's argumentation, these cycles should mainly be found in the most competitive markets, whereas less competitive markets should rather be characterized by sticky or constant prices. Consequently, if Edgeworth price cycles are the correct explanation for a perceived asymmetric price adjustment, policy makers could make a big mistake by (only) punishing firms who compete in markets where such a pricing behavior is apparent.

The pattern Eckert describes in his model has nothing to do with 'true' asymmetric pricing, which is the sole focus of my thesis. However, it is straightforward to

[33]If marginal cost is constant, only the lowest price in the cycle will be a function of cost.

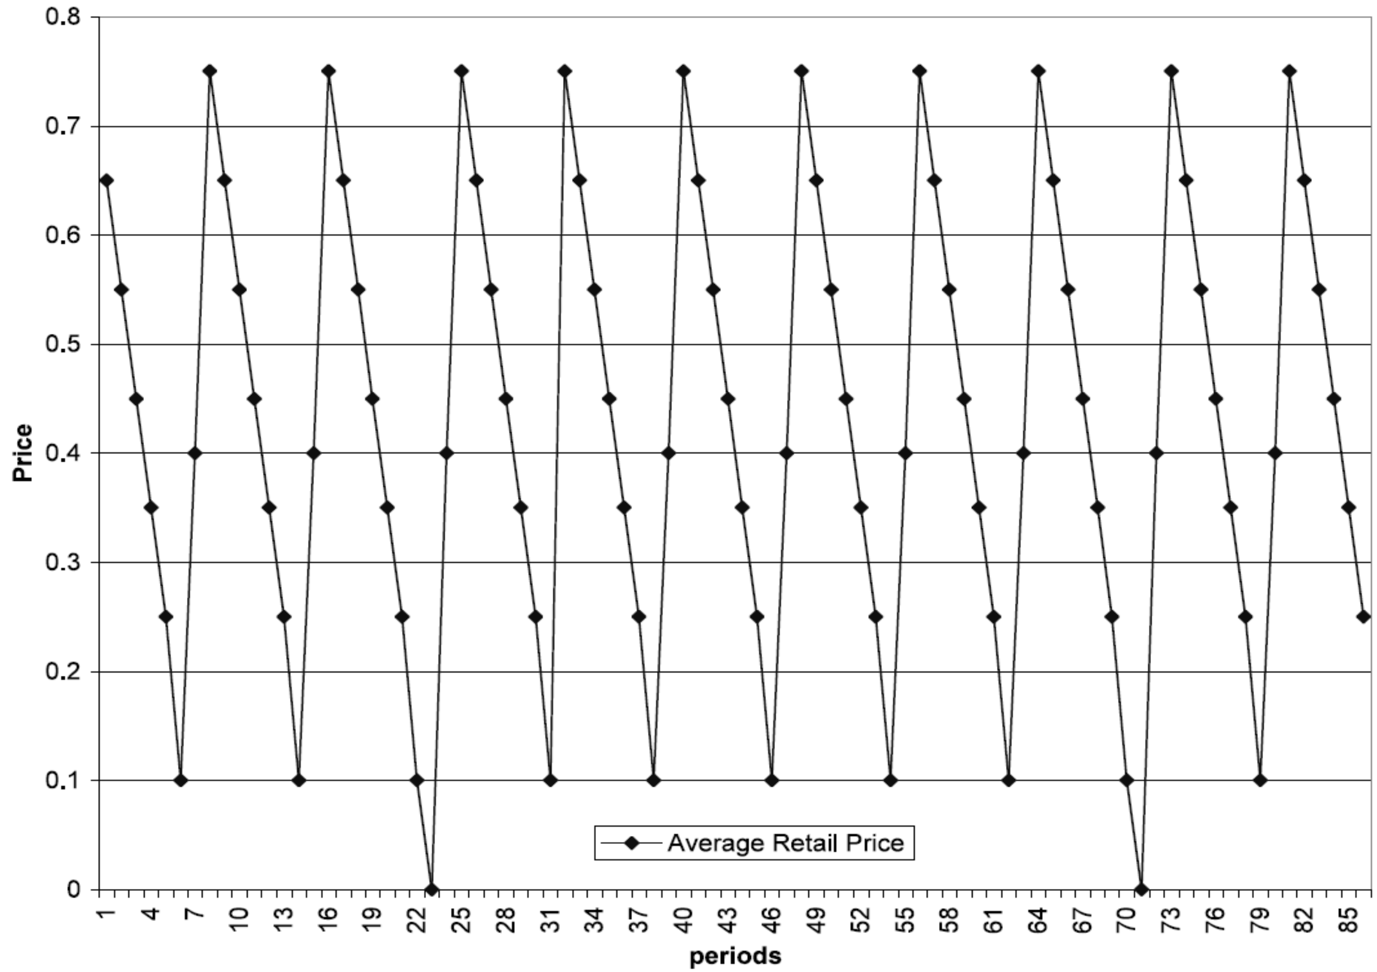

Figure 5.1 Example for an Edgeworth price cycle with constant marginal cost. Source: Eckert (2002, p. 63)

see that an extension of his model could lead to rockets and feathers in the traditional sense. If marginal cost is not assumed to be constant, but can change with a sufficiently low probability, Eckert argues that firms will not alter the strategy outlined above in equilibrium. If costs change unexpectedly while firms are in the undercutting regime (which they are most of the time), it seems reasonable to think that the undercutting regime would continue until the new marginal cost is reached, given that prices are still above cost. This is because firms try to steal market shares in the undercutting regime, which shouldn't be affected by lower or higher marginal cost. If marginal cost rises to such an extent that margins become negative, however, firms are expected to immediately increase prices to the (new) highest point of the cycle.
In sum, asymmetric price adjustment will be found: negative cost shocks are transmitted slowly to selling prices (as the undercutting regime is simply prolonged) whereas positive cost shocks *can* have an immediate effect on prices if margins become negative. But even if not, the undercutting regime lasts shorter, meaning that positive cost shocks will be transmitted more quickly to output prices than negative ones. This implies asymmetric pricing in the traditional sense.

In summary, I think that Edgeworth price cycles à la Eckert (2002) might turn out to be one of the most important sources of asymmetric pricing in highly competitive markets. If such cycles were found to be common in many competitive markets, the theory of asymmetric pricing would have to be reconsidered. Then, other explanations like implicit collusion or costly consumer search would lose much of their appeal.
On the other hand, if cost shocks were usually found to be transmitted in the classic sense of rockets and feathers, with only few markets exhibiting a pattern of Edgeworth price cycles, new theories would be needed to pin down the reasons that

discern classical asymmetric pricing from asymmetric pricing caused by Edgeworth cycles. In any case, a lot of further empirical and theoretical research will be necessary to give a better understanding of the issue.

6 Summary and Conclusion

In this thesis, the phenomenon of asymmetric adjustment of output prices to input price shocks was examined in detail. After pointing out why asymmetric pricing is a relevant topic in contemporaneous economic research in the introduction, I started out by providing a definition and distinction between the various types of the so called pattern of rockets and feathers.
In the main section of this work, I proceeded to develop a model of oligopolistic coordination that resulted in asymmetric price adjustment. In particular, my intention was to bridge the gap between the numerous research papers that reference collusion as one of the main sources of asymmetric pricing and the absence of *any* formal model that is actually able to derive asymmetric pricing caused by collusion.
Finally, I summarized several other models that lead to rockets and feathers in the last two sections of this thesis. While consumer search models are currently the dominating category, other interesting ideas have been developed that might renew our understanding of asymmetric pricing.

Overall, one can conclude that there are numerous mechanisms that could possibly result in asymmetric pricing. As of now, it is still unclear whether rockets and feathers is a product of collusion, imperfect consumer search, miscellaneous sources (like, for example, lags in adjustment of production or Edgeworth price cycles) or a combination of all or some of those.
In this thesis, I was able to prove that a simple dynamic oligopoly model can lead to asymmetric pricing. If overall demand is random and enough consumers in a market know the prices of all competing firms, firms might find it optimal to collude on the price that was charged before a negative cost shock if costs drop,

but increase prices immediately if costs increase. Due to the random component of demand, collusion can slowly break down after a negative cost shock, leading to asymmetric pricing.

The model I contributed shares one feature with most other models of asymmetric pricing: a higher persistence of costs leads to a more severe adjustment asymmetry of prices. Also, more competition in a market leads to a reduced likelihood of collusion and thus asymmetric pricing, which is a plausible result. On the other hand, the property that a higher fraction of informed consumers strengthens asymmetric price transmission is not found in any other model I examined. While this prediction of the trigger sales model may seem problematic at first, it is not surprising in its context: if the fraction of informed consumers is small, firms have essentially no chance to punish deviating firms under collusion. Put differently, the trigger sales model provides a benchmark to distinguish asymmetric pricing caused by collusion from asymmetric pricing caused by imperfect consumer search and (at least some) other mechanisms. Asymmetric pricing under collusion should only be found if consumers are very well informed about the prices in a market. If consumers have limited knowledge about them, collusion becomes a less likely explanation for the phenomenon of rockets and feathers.

Bibliography

Robert W. Bacon. Rockets and feathers: the asymmetric speed of adjustment of uk retail gasoline prices to cost changes. *Energy Economics*, 13(3):211–218, July 1991.

Laurence Ball and N. Gregory Mankiw. Asymmetric price adjustment and economic fluctuations. *Economic Journal*, 104(423):247–61, March 1994.

Severin Borenstein, A Colin Cameron, and Richard Gilbert. Do gasoline prices respond asymmetrically to crude oil price changes? *The Quarterly Journal of Economics*, 112(1):305–39, February 1997.

Luìs Cabral and Arthur Fishman. Business as usual: A consumer search theory of sticky prices and asymmetric price adjustment. New York University and CEPR (Luìs Cabral); Bar Ilan University (Arthur Fishman), August 2008.

Andrew Eckert. Retail price cycles and response asymmetry. *Canadian Journal of Economics*, 35(1):52–77, February 2002.

Ronald Johnson. Search costs, lags and prices at the pump. *Review of Industrial Organization*, 20(1):33–50, February 2002.

Jeffrey D. Karrenbrock. The behavior of retail gasoline prices: symmetric or not? *Review*, (Jul):19–29, 1991.

Matthew Lewis. Asymmetric price adjustment and consumer search: An examination of the retail gasoline market. The Ohio State University, Department of Economics, September 2009.

Jochen Meyer and Stephan Cramon-Taubadel. Asymmetric price transmission: A survey. *Journal of Agricultural Economics*, 55(3):581–611, 2004.

Sam Peltzman. Prices rise faster than they fall. *Journal of Political Economy*, 108 (3):466–502, June 2000.

Mariano Tappata. Rockets and feathers: Understanding asymmetric pricing. *RAND Journal of Economics*, 40(4):673–687, 2009.

R Development Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2009. URL `http://www.R-project.org`.

Jean Tirole. *The Theory of Industrial Organization*. Cambridge, MA: MIT Press, 1988.

Hal R. Varian. A model of sales. *American Economic Review*, 70(4):651–59, September 1980.

Wikipedia. Asymmetric price transmission — wikipedia, the free encyclopedia, 2009. URL `http://en.wikipedia.org/w/index.php?title=Asymmetric_price_transmission&oldid=305200103`. [Online; accessed 11-January-2010].

Wikipedia. Folk theorem (game theory) — wikipedia, the free encyclopedia, 2010. URL `http://en.wikipedia.org/w/index.php?title=Folk_theorem_(game_theory)&oldid=336417180`. [Online; accessed 11-March-2010].

Huanxing Yang and Lixin Ye. Search with learning: understanding asymmetric price adjustments. *RAND Journal of Economics*, 39(2):547–564, 2008.

List of Figures

2.1 Time, amount and pattern asymmetry 7
2.2 Amount asymmetry with negative asymmetric price transmission . 10
3.1 Depiction of collusive trigger sales equilibria 28
3.2 Probability distribution function of the number of periods until collusion collapses in a simple trigger sales equilibrium 33
3.3 A simulation of average market prices in a simple trigger sales equilibrium 34
3.4 Probability mass function of the number of submarkets who break collusion over time in a trigger sales equilibrium with multiple separated submarkets 38
3.5 A simulation of average market prices in a trigger sales equilibrium with multiple separated submarkets 40
3.6 A simulation of how collusive submarkets collapse 41
4.1 Example for a unique equilibrium in Tappata's model 70
5.1 Example for an Edgeworth price cycle with constant marginal cost . 90

Appendix A

A.1 Abstract

This thesis studies the phenomenon of asymmetric pricing, according to which output prices tend to adapt quicker to input cost increases than to input cost decreases. Although the general public and government authorities tend to attribute this pricing behavior to an abuse of market power, no formal model capable of generating asymmetric pricing through collusion has been developed. The main contribution of this thesis is to provide such a model. In my model, positive cost shocks lead firms' margins to be squeezed, implying an immediate response of output prices. In contrast, firms try to coordinate their prices after negative cost shocks because the old output price provides a natural focal point for collusion. However, overall market demand is random and unobservable to firms. As a consequence, they confuse random demand shocks with demand shocks caused by rival firms undercutting, leading collusion to eventually break down. Overall, asymmetric pricing is the result.

Several other models of asymmetric pricing are presented in this thesis. While consumer search models typically imply a less pronounced form of asymmetric pricing if many consumers are informed (i.e., observe all prices in the market), my model predicts the opposite. The reason is that collusion under random and unobservable demand can only be maintained if firms can successfully punish a deviating firm. This is unlikely to happen if deviation has little effect, which is the case if there are few informed consumers.

A.2 Abstract (German)

Diese Diplomarbeit befasst sich mit dem Phänomen der asymmetrischen Preistransmission, welches eine schnellere Reaktion von Verkaufspreisen auf positive Kostenschocks im Vergleich zu negativen Kostenschocks impliziert. Obwohl die Öffentlichkeit sowie Regierungsstellen asymmetrische Preistransmission häufig mit Marktmissbrauch gleichsetzen, wurde bisher noch kein formales Modell entwickelt, welches asymmetrische Preistransmission durch Kollusion erklären kann. Der wichtigste Beitrag dieser Arbeit is daher, ein solches Modell zu entwickeln. In meinem Modell führen positive Kostenschocks zu einer unmittelbaren Veränderung der Verkaufspreise, da die Verkaufsmarge sonst negativ würde. Im Gegensatz dazu veranlassen negative Kostenshocks die Firmen, ihre Preise auf dem alten Verkaufspreis zu koordinieren, da dieser Preis einen natürlichen Fokuspunkt für Kollusion darstellt. Da aber in meinem Modell die Gesamtnachfrage zufällig und unbeobachtbar ist, verwechseln die Firmen zufällige Nachfrageschocks mit Nachfrageschocks, die durch unterbietende Firmen enstehen. Dies führt dazu, dass Kollusion nach einer gewissen Anzahl an Perioden zusammenbricht. Insgesamt kommt es zu asymmetrischer Preistransmission.

Diverse andere Modelle asymmetrischer Preistransmission werden ebenfalls in dieser Diplomarbeit präsentiert. Während Consumer-Search Modelle typischerweise eine schwächere Form von asymmetrischer Preistransmission bei einer geringen Anzahl von informierten Konsumenten (d.h. Konsumenten, die sämtliche Preise im Markt beobachten) implizieren, sagt mein Modell das Gegenteil voraus. Der Grund hierfür ist, dass Kollusion unter zufälliger und unbeobachtbarer Gesamtnachfrage nur dann aufrecht erhalten werden kann, wenn Firmen erfolgreich eine abweichende Firma bestrafen können. Dies ist aber unwahrscheinlich, wenn Abweichung einen sehr kleinen Effekt hat, wie das bei einer geringen Anzahl von informierten Konsumenten der Fall ist.

A.3 Curriculum Vitae

Name:

Martin Obradovits

Date of birth:

December 13th, 1985

Place of birth:

Vienna, Austria

Education:

October 2005 to April 2010: Studies of Economics at the University of Vienna

Since October 2008: Studies of Statistics at the University of Vienna

June 2004: Matura with honours, BRG Rahlgasse, 1060 Vienna