The London School of Economics and Political Science

Product Differentiation, Uncertainty and Price Coordination in Oligopoly

Thesis submitted by

Michael Alexander Raith

to the

University of London

for the degree of

Doctor of Philosophy

January 1996

To my parents

Abstract

This thesis consists of three self-contained analyses of models with price-setting firms. It explores the relationships between different sources of market imperfection that may be present simultaneously: product differentiation, imperfect information and collusive pricing.

Chapter 2 analyses the circumstances under which oligopolists have an incentive to exchange private information on unknown demand or cost parameters. It presents general model which encompasses virtually all models in the current literature on information sharing as special cases. Within this unifying framework it is shown that in contrast to the apparent inconclusiveness of previous results, some simple principles determining the incentives for firms to share information can be obtained. Existing results are generalised, some previous interpretations questioned and new explanations offered.

Chapter 3 addresses the question of how price setting between firms in a spatial retail market is affected if the relevant consumers commute between their home and their workplace and try to combine shopping with commuting. It is shown within a specific model that for small commuting distances, an increase in commuting leads to a decrease of equilibrium prices, since due to a reduction of effective travel costs the firms' products become better substitutes. Under quite general conditions, however, larger or dispersed commuting distances lead to the nonexistence of a price equilibrium.

Chapter 4 analyses the question how product differentiation affects the scope for oligopolists to collude on prices. It suggests a precise theoretical foundation for the conventional view that heterogeneity is a factor hindering collusion, a view which has been challenged in recent theoretical work. It is argued that, in a world of uncertainty, an increase in the heterogeneity of products leads to a decrease in the correlation of the firms' demand shocks. With imperfect monitoring, this makes collusion more difficult to sustain, as discriminating between random demand shocks and deviations from the cartel strategy becomes more difficult.

Contents

List of Figures

Acknowledgements

In the preparation of this thesis I have, over the years, received much encouragement and advice by mentors, colleagues and friends.

My greatest and most obvious debt is to my supervisor Professor John Sutton. From my vague and unfocused ideas to the completed thesis, he has continuously, encouragingly and with much patience guided me through all stages of my work. Through his Economics of Industry graduate course, the "Self-Help Group", and countless meetings, I have learned tremendously, and his ideas and approach have shaped my views on economics.

I would also like to thank Professor Urs Schweizer, who stimulated my interest in Industrial Economics and, as the supervisor of my diploma thesis in Bonn, encouraged the work that was the basis of Chapter 2. In addition, he has provided much support both personally and as coordinator of the European Doctoral Programme in Quantitative Economics.

Professor Patrick Bolton's comments and ideas have always been a source of inspiration and challenge. Whenever I looked for advice, he has motivated and supported me, for which I am very grateful.

Thanks are also due to Professor Xavier Vives for his active interest in my work, for giving me the opportunity to present it at different occasions, and for insightful suggestions.

The work on this thesis has benefited from the help of numerous other people. Special thanks go to Guido Friebel, who over the years has provided me with detailed and critical comments on various versions of all parts of this thesis.

Chapter 2 originates from my diploma thesis submitted to the University of Bonn in December 1991. Different versions of it have benefited from comments by Frank Bickenbach, Rembert Birkfeld, Clemens Esser, Benny Moldovanu, Georg Nöldecke, Avner Shaked and Mark Spoerer. Substantial improvements were suggested by a referee of the *Journal of Economic Theory.*

The work on Chapter 3 started with discussions with Alison Hole and Reinout

Koopmans. Michael Hardt provided helpful comments on this chapter, and many improvements of the exposition, as well as an extension, were suggested by referees of the *International Journal of Industrial Organization.*

Chapter 4 has had a somewhat turbulent history since the drafting of the first version in February 1995. In the course of its evolution, comments by Sougata Poddar, George Symeonidis, seminar participants at Studienzentrum Gerzensee (especially Preston McAfee), and Tore Ellingsen proved to be particularly influential. Godfrey Keller carefully read the final version and suggested many refinements.

Moreover, I have received much useful feedback on what are now Chapters 2 and 4 from seminar participants at LSE, the EDP Jamborees, and various summer schools, workshops and conferences.

Most of the work on this thesis was carried out at the Suntory-Toyota International Centres for Economics and Related Disciplines at LSE with its stimulating atmosphere and its excellent research facilities. Thanks are due to Luba Mumford and Nora Lippincott, who have always been very helpful, and to everybody else in the Economics of Industry Group. The final work on the thesis was carried out at the European Centre for Advanced Research in Economics in Brussels.

I am grateful to several organizations for sponsoring my participation at conferences and workshops: the Centre for Economic Policy Research, the European Association for Research in Industrial Economics, LSE, the Review of Economic Studies, Studienzentrum Gerzensee, and Tel Aviv University.

Finally, I wish to express my thanks for the generous financial assistance I have received, in different years, by the Deutsche Forschungsgemeinschaft through SFB 303, the German Academic Exchange Service (DAAD), the Suntory-Toyota International Centres for Economics and Related Disciplines, and the European Commission through its Human Capital and Mobility Programme.

Chapter 1

Introduction and Summary

In the study of imperfect competition among firms, most of the papers in the theoretical literature have focused on just one of four sources of imperfection: competition in quantities à la Cournot, product differentiation, imperfect information and collusion. The roles of these aspects of imperfection are quite well understood where they appear in isolation. It is much less well understood, however, how markets operate in which two or more of these factors are present simultaneously. The purpose of this thesis is to explore some of these relationships between different sources of imperfection.

The thesis comprises three self-contained game-theoretic analyses of oligopolistic interaction among firms. While these essays are concerned with quite different questions, a common characteristic of all three is the study of *price setting* firms in industries with *differentiated products.* This reflects what seems to be a general tendency in recent theoretical contributions in industrial economics: the departure from the use of the Cournot model of quantity competition as the standard "workhorse oligopoly theory" (Shapiro 1989, p. 346).

The popularity of the Cournot model was not so much based on the belief that firms typically compete in quantities. Rather, it was seen as an analytically tractable model which led to apparently more realistic predictions than the homogeneous goods Bertrand model and therefore appeared as a suitable metaphor for imperfect price competition. The Cournot model gained further credibility through its depiction as a 'capacity choice' model by Kreps and Scheinkman (1983).

It has turned out, however, that for entire classes of games - usually two-stage games, an example of which will also be studied in this thesis - quantity and price competition can lead to opposite results. This reflects, in part, the fact that the firms' reaction functions carry opposite signs in the two models. This fundamental difference raises doubts as to whether the Cournot model really is a suitable substitute for explicit models of price-setting.

On the other hand, if standard game-theoretic concepts are used, any meaningful study of models with price-setting firms is possible only in the context of product differentiation, since otherwise the unrealistic predictions of the homogeneous-goods Bertrand model arise. The theory of product differentiation, now, has made major advances over the years. Today, theorists have a large repertoire of models of (horizontal or vertical) product differentiation to choose from, and there exists a theory which unifies these different approaches (Anderson, de Palma and Thisse 1992). Thus, this development of the theory of product differentiation makes the study of price competition in particular applications less intractible than it used to be.

"Price coordination" refers to both noncooperative and collusive pricing among firms. One issue that has been explored is the way in which horizontal product differentiation leads to higher prices in a noncooperative (Nash) equilibrium, due to a relaxation of price competition.

Chapter 3 studies what might be seen as a counterexample to this general relationship. Here, product differentiation of a very special kind is considered. Consumers in a spatial retail market are assumed to commute between their home and their workplace and try, where possible, to combine commuting with shopping. Now, an increase in the (average) distance over which consumers commute makes the retail firms' product(s) better substitutes in the sense that the travel costs associated with making a purchase decrease. Hence, the degree of product differentiation between the firms is, other things being equal, a function of commuting distance, and the above reasoning suggests that as a result of an increased commuting distance, the retail market becomes more competitive. The analysis of equilibrium prices in this kind of market, however, reveals some deviations from the standard relationship between product differentiation and noncooperative pricing.

Another issue is addressed in Chapter 4, viz. the relationship between product differentiation and the scope for collusive pricing. This issue has recently received much attention by theorists, and many have questioned the validity of the conventional wisdom that heterogeneity of products is a factor hindering collusive conduct. Chapter 4 of this thesis offers a new approach to the analysis of this relationship.

Uncertainty plays an important role in Chapters 2 and 4, both of which originate from questions that have been the subject of substantial discussion in the area of competition policy, more specifically, questions related to the assessment of horizontal agreements between firms. It used to be believed that uncertainty plays a double role in its impact on the firms' conduct. On the one hand, where firms compete, it seems desirable, both from the firms' and from a social point of view, to have the greatest possible market transparency, since uncertainty causes inefficient frictions. On the other hand, sufficient transparency - both with respect to firms' actions and to the environment - is also seen as a necessary condition for firms to be able to collude on prices.

Chapters 2 and 4 address issues related to this line of reasoning. One of these is the question how the firms' conduct, and social welfare, are affected by the exchange of information between firms in the same industry. This question was first raised when information sharing agreements in various industries and countries came under the scrutiny of antitrust authorities, on the grounds that they allegedly facilitated collusive behaviour. Motivated by this, a large theoretical literature emerged during the 1980s, which has come to focus on the question whether information exchange is profitable for firms if they do not collude. Perhaps surprisingly, the first results suggested that this is not he case. This led to the argument that the observation of information-sharing agreements should be regarded as prima-facie evidence of collusion (Clarke 1983). Chapter 2 is concerned with this strand of literature. It presents a general framework in which the firms' incentives to exchange information can be studied.

The second application of the above reasoning relates to the conventional wisdom that heterogeneity of products impedes collusive conduct because it implies a situation of increased "complexity" or uncertainty. Yet, while the link between uncertainty and cartel stability seems fairly uncontroversial, it is much more difficult to pin down analytically how product differentiation is related to "complexity", since in standard models, product differentiation affects the demand elasticities but not the degree of uncertainty or complexity. In fact, theorists have in recent years appealed to their analyses of deterministic models in challenging the conventional view. Chapter 4 constitutes an attempt to provide a theoretical foundation for the conventional view.

In the following, I shall summarise the results of the chapters that follow.

Chapter 2[1] analyses the circumstances under which oligopolists have an incentive to share private information about stochastic demand or stochastic costs. Pioneered by the works of Novshek and Sonnenschein (1982), Clarke (1983) and Vives (1984), a large literature exploring this question emerged during the past decade. While the models analysed vary along several dimensions, their basic structure is the same. It has turned out that the results of the models depend delicately on the specific assumptions used. Hence, little is known about which main forces drive the incentives of firms to share private information.

In this chapter I develop and analyse a general model of information sharing in oligopoly. The model is constructed so as to encompass virtually all models in the current literature as special cases.

The structure of the model is the same as that used in the literature: (i) In an n-firm oligopoly with differentiated goods, firms face either a stochastic intercept of a linear demand function or a stochastic marginal cost, which can be different for each firm. The deviation of the vector of demand intercepts or costs from its mean is unknown to the firms. (ii) Instead, each firm receives a private signal with information about the true state of nature. For example, firms might receive noisy signals about the intercept of a common demand function; or they might know their own costs exactly, but not the costs of the rival firms. (iii) Private information can be exchanged. Following the literature, I assume that firms commit themselves either to reveal their private information to other firms, or to keep it private, *before* receiving any private information. (iv) In the last stage, the "oligopoly game", firms noncooperatively set prices or quantities so as

[1] forthcoming in *Journal of Economic Theory*.

to maximise expected profits conditional on the available private and revealed information.

Two different approaches are used to analyse the revelation behaviour of firms. In the simpler case, we determine the conditions under which industry-wide contracts on information sharing are profitable, by comparing the expected equilibrium profits with and without information sharing. Alternatively, we assume that firms decide on their revelation behaviour simultaneously and independently, thus allowing for asymmetric revelation decisions.

The analysis of the model not only leads to results more general than previous ones; more importantly, a small number of forces driving the incentives to share information in most types of models can be identified. I argue that previous interpretations of information sharing models are not always consistent with the formal analyses.

Among other results, it is shown that for Cournot markets, and for Bertrand markets with demand uncertainty, there are some simple general results underlying almost all previous results: with perfect (i.e. noiseless) signals or uncorrelated demands/costs, or with a common value and strategic complements, complete information pooling is an equilibrium of the two-stage game (which is efficient from the viewpoint of the firms), regardless of all other parameters. With a common value and strategic substitutes, no pooling is the equilibrium solution. This solution is efficient in Cournot markets with homogeneous goods and inefficient if the degree of product differentiation is large. It is shown that the profitability of information sharing in most models with cost uncertainty is driven merely by the assumption that firms know their own costs with certainty. This suggests that certain interpretations of results that have been suggested in the literature are invalid.

For these types of models, I propose a new explanation for the incentives to reveal private information which rests on two principles: (i) Letting the rivals acquire a better knowledge of their respective profit functions leads to a higher correlation of strategies, the profitability of which is determined by the slope of the reaction curves. (ii) Letting rivals acquire better knowledge of one's own profit function always increases the firm's own expected profits by inducing a change of the correlation of strategies in the direction which is profitable. The incentive to reveal information is then determined by the sum of these two effects.

Chapter 3[2] is concerned with the question how the interaction between firms in a spatial retail market is affected if the relevant customers are assumed to commute between their home and their workplace.

A salient feature of most models of spatial competition is the assumption that consumers, like firms, can appropriately be characterized by a single point in address space. Moreover, the cost of travelling from a consumer's (home) location to a firm (a retail store) is fully attributed to the net price the consumer has to pay. While this may be a reasonable assumption in models where the spatial dimension refers to abstract characteristics, it can be quite unrealistic in the original context of geographical distribution. That is, for a large variety of retail markets, the relevant consumers are more appropriately described as living at one location and working at another. Commuting between these locations occurs regardless of any purchases made, and consumers will in general try to combine shopping with commuting in order to save travel costs. In this case, the net travelling distance associated with the purchase of a good should not include any travelling along the commuting route.

Then some of the questions that arise are the following: How does commuting affect prices for any given locations of firms? How does commuting affect location choice, both in the case of fixed prices and in the case of anticipated price competition? Can asymmetric equilibria emerge in this context? A brief discussion in the Introduction of Chapter 3 reports some of the results that can be obtained. It turns out that the main problem is that in various settings, the introduction of commuting into an otherwise standard model is likely to give rise to the nonexistence of an equilibrium in prices or quantities.

[2] forthcoming in *International Journal of Industrial Organization*.

The case of price competition for given locations is subsequently analysed in some detail using a model due to Claycombe (1991). In Claycombe's (1991) model, firms and consumers are located on an infinite line. Firms are spaced at an equal distance, and consumers are uniformly distributed along the line. Each consumer commutes to a location at a constant distance from his home location. Claycombe argues that commuting makes the market inherently more competitive. In particular, he arrives at the result that if the commuting distance exceeds the distance between the firms, Nash price setting behaviour essentially leads to a competitive outcome, whereas for smaller commuting distances, commuting does not matter very much.

In a subsequent study, Claycombe and Mahan (1993) interpret this theoretical work as implying that if a population consists of both commuting and noncommuting consumers, the equilibrium retail price depends negatively on the commuting distance of the commuters and positively on the fraction of noncommuting consumers. They go on to estimate retail prices for beef in different U.S. cities, including commuting characteristics as explanatory variables. Their results seem to offer support to the predictions. It turns out that the proportion of consumers using mass transit or car pools has a significant positive effect on prices, whereas the average commuting distance has a marginally significant negative impact.

In this essay, I apply a rigorous game-theoretic analysis to a simplified version of the original Claycombe model. It turns out that the results differ considerably from those obtained by Claycombe and at the same time are broadly consistent with the predictions and empirical results of Claycombe and Mahan (1993).

For small commuting distances, prices in a symmetric equilibrium depend continuously and negatively on the commuting distance and positively on the proportion of noncommuting consumers. For larger commuting distances, however, a symmetric price equilibrium in pure strategies in general does not exist. The reason is quite interesting. With an increase in commuting distance, the market becomes, in a sense, more competitive, and profits decrease. At some point, however, firms have an incentive to cease competing for the marginal consumers and instead *increase* their price in order to extract profits from local consumers over whom they have some monopoly power. Thus, it is the coexistence of two groups of consumers that quite generally leads to the nonexistence of an equilibrium in such games of price competition. Consumers that costlessly pass two or more firms on their commuting route draw the firms into intense price competition, the result of which would be marginal cost pricing. On the other hand, the existence of consumers that pass only one firm implies that firms can always secure supramarginal profits. Thus, while nonexistence results are fairly common in spatial models, the reason for the breakdown of equilibrium in our model is quite different from the reasons for breakdown in some other models. Only for large commuting distances, viz. at least twice the distance between firms, and only if all consumers commute, does perfect competition prevail (since in that case every consumer passes at least two firms on his or her commuting route).

Chapter 4 analyses the effect of product differentiation on the stability of collusion. According to the conventional view, product heterogeneity entails a situation of "higher complexity" than prevails with homogeneous products. This complexity limits the scope for collusion because it makes both an agreement on a collusive strategy and its subsequent enforecment more difficult. An analytical formulation of this argument, however, has not been available.

This view has been challenged in recent theoretical work. Theorists have argued that since for differentiated goods both the gain from deviating and the benefit of colluding (equivalently, the severity of punishment) are small compared to the case when goods are homogeneous, little can be said *a priori* on how product differentiation affects cartel stability. In fact, the analysis of various different models has even suggested a positive relationship between differentiation and the stability of collusion.

In this chapter, I introduce the idea that, in a world of uncertainty, an increase in the heterogeneity of products leads to a decrease in the correlation of the firms' demand shocks. If firms sell homogeneous products, they attract the same groups of customers, and shocks on the demand side will affect firms in the same way. In contrast, with differentiated goods, firms attract different customers, and hence shocks on the demand side affect firms differently, i.e. the demand functions are less correlated. This is illustrated with a variant of the Hotelling model with taste heterogeneity due to de Palma et al. (1985). Here, we show how the introduction of demand uncertainty leads to stochastic demand functions for the two firms such that their correlation coefficient is is an increasing function of the degree of substitutability. Both in terms of the result and the underlying economic intuition, the model precisely captures the idea suggested above.

Now, if each firms cannot observe the other's price but has to infer from observed demand whether another firm has deviated, the correlation effect makes collusion more difficult to sustain when goods are differentiated (as discriminating between random demand shocks and deviations from the cartel strategy becomes more difficult). This effect is illustrated with two duopoly models, in which for simplification the correlation effect is introduced in a rather *ad hoc* way. These are models in which price-setting firms use strategies that are optimal within a certain class of generalised Green-Porter (1984)-type trigger strategies. The first model is a Hotelling-type model, the second a model with a demand system derived from a 'representative consumer' utility function.

While the properties of these models are quite different, they generate similar predictions with regard to collusion: given a sufficiently high discount factor and a sufficient level of demand uncertainty, collusion becomes less sustainable as products become more differentiated. Moreover, if differentiation exceeds a critical level, collusion may not be profitable at all. These results stand in sharp contrast to those in the recent theoretical literature. At the same time, this theory provides a simple analytical foundation for the traditional view that heterogeneity limits the scope for collusion.

From a technical point of view, the analysis synthesises two strands of literature. One strand deals with deterministic models. Here, if a firm deviates from a collusive strategy, it seeks to maximise its profit per period, since by assumption any deviation is detected and it precipitates an immediate retaliation. The other strand of literature, starting with Green and Porter (1984), has analysed the sustainability of collusion in an imperfect-monitoring framework by looking at the profitability of *marginal* deviations from a collusive strategy. In Chapter 4, I emphasize the existence of a fundamental nonconcavity in a firm's payoff function, which implies that for the analysis of collusion both large discrete deviations and marginal deviations must be considered simultaneously.

Chapter 2

A General Model of Information Sharing in Oligopoly

2.1 Introduction

Theoretical research on information sharing in oligopoly was pioneered by Novshek and Sonnenschein (1982), Clarke (1983) and Vives (1984). Since then, numerous contributions on this topic appeared. While the models analysed vary along several dimensions, their basic structure is the same. According to the received view on the current state of this field, there is no general theory regarding the incentives of firms to share private information; rather, the results of the models depend delicately on the specific assumptions.

In this chapter I analyse a general model of information sharing in oligopoly. The model is constructed so as to encompass virtually all models in the current literature as special cases, resulting from appropriate specification of the parameters. The analysis of the model not only leads to results more general than previous ones; more importantly, a small number of forces which drive the incentives to share information in most types of models can be identified. It is argued that previous interpretations of information sharing models are not always consistent with the formal analyses, and I suggest a new explanation for the incentives to reveal information.

The structure of our model is the same as that used in the literature: (i) In an n-firm oligopoly with differentiated goods, firms face either a stochastic intercept of a linear demand function or a stochastic marginal cost, which can be different for each firm. The deviation of the vector of demand intercepts/costs from its mean, the "State of Nature", is unknown to the firms. (ii) Instead, each firm receives a private signal with information about the true State of Nature. For example, firms might receive noisy signals about the intercept of a common demand function; or they might know their own costs exactly, but not the costs of the rival firms. (iii) Private information can be exchanged, where it is assumed that firms commit themselves either to reveal their private information to other firms or to keep it private *before* receiving any private information. (iv) In the last stage, the "oligopoly game", firms noncooperatively set prices or quantities so as to maximise expected profits, conditional on the available private and revealed information.

Following the literature, I use two different approaches to analyse the revelation behaviour of firms. In the simpler case, the conditions under which industry-wide contracts on information sharing are profitable are determined by comparing the expected equilibrium profits with and without information sharing. Alternatively, we assume that firms decide on their revelation behaviour simultaneously and independently, thus allowing for asymmetric revelation decisions.

How, then, do prices or quantities and expected profits with and without information sharing depend on the characteristics of the market, and how does this affect the incentives for firms to exchange information in the first place?

I will not review in detail the various contributions addressing these questions; a brief survey can be found in Vives (1990). It has been noted by Vives (1990) and others that the results concerning the incentives to share information seem to depend sensitively on the specific assumptions of the model: a change from Cournot to Bertrand competition, from substitutes to complements, from demand to cost uncertainty, or from a common value to "private values", referring to an n-dimensional State of Nature for n firms, may lead to completely different outcomes. More disturbingly, apparently similar models often lead to contrasting results. Three points shall illustrate that the current literature cannot satisfactorily explain the diversity of results, or worse, that only little seems to be known about the forces driving each particular result.

(i) According to the received view, there are two main effects of information sharing from the viewpoint of the firms (excluding collusion in the price/quantity setting stage). On the one hand, each firm is better informed about the prevailing market conditions, which is presumably profitable. On the other hand, the homogenization of information among firms leads to a change in the correlation of the strategies. An increase of the correlation in turn is profitable for Bertrand competition but not for Cournot competition. The overall profitability is then determined by the sum of these two effects.

Except for some special cases, however, this well-known reasoning is either inapplicable or flawed. First, if a firm is perfectly informed about its own cost, it is in general not true that it benefits from obtaining information about its rivals as well (Fried 1984, Sakai 1985). Second, the change in the correlation of strategies is itself endogeneous and not easily predicted. I show that contrary to what is sometimes believed it is in general not true that information sharing always leads to an increase of the correlation, or that the correlation increases with a common value and decreases with private values. Finally, it also shown that even the relationship mentioned above between a change in the correlation of strategies and its effect on expected profits (depending on the slopes of the reaction curves) is not quite as generally valid as is usually assumed.

(ii) Vives (1984), Gal-Or (1985), and Li (1985) have shown that in a Cournot oligopoly with homogeneous goods and demand uncertainty , firms do not share information in the equilibrium of the two-stage game described above. In contrast, Fried (1984), Li (1985), and Shapiro (1986) have shown that in a Cournot market with uncertainty about private costs, firms completely reveal information in the equilibrium. This contrast has been attributed to the difference between a common value, e.g. the intercept of a common demand function, and private values, e.g., different marginal costs for the firms. But the results for private values in many models hold even if the correlation of marginal costs approaches unity, although economically, this situation is equivalent to a model with a common value. Therefore, this interpretation is inconsistent with the models as it suggests a discontinuity of profits in the underlying parameters which one would not expect in this class of models.

(iii) Vives (1984) shows that in a duopoly with differentiated products and demand uncertainty, a change from substitutes to complements or from Cournot to Bertrand yields opposite results regarding the incentives to share information. This may be attributed to a change in the slope of the reaction curves. However, in the private-values, cost-uncertainty model of Gal-Or (1986) there is only a difference between Cournot and Bertrand but not between substitutes and complements. Finally, in Sakai's (1986) model firms always share information, regardless of whether they set prices or quantities, or whether the goods are substitutes or complements. Hence from these results, very little can be concluded about the role of the type of competition and the characteristics of goods.

The results of the analysis in this chapter imply that these interpretative problems can all be resolved:

1. I argue (prior to the formal analysis) that the distinctive characteristic of most so-called private-value models is not the independence of (say) costs, but the fact that firms are perfectly informed about their own cost. I therefore introduce a new distinction between independent-value models and what I label "perfect-signal" models.

2. The correlation of strategies increases with information sharing in the case of a common value. With independent values or perfect signals, however, the direction of change depends on the slope of the reaction curves.

3.(a) For Cournot markets, and for Bertrand markets with demand uncertainty, there are some simple general results underlying almost all previous results: with perfect signals or uncorrelated demands/costs, or with a common value and strategic complements, complete information pooling is an equilibrium of the two-stage game (which is efficient from the viewpoint of the firms), regardless of all other parameters. With a common value and strategic substitutes, no pooling is the equilibrium solution. This solution is efficient in Cournot markets with homogeneous goods and inefficient if the degree of product differentiation is large. It is shown that the profitability of information sharing in most models with cost uncertainty is driven merely by the assumption that firms know their own costs with certainty, which refutes previous interpretations attributing these results to other factors.

(b) For the remaining case, Bertrand markets with cost uncertainty, the incentives are rather ambiguous. It turns out that results derived for duopoly models (Gal-Or 1986) may be reversed in the case of many firms. More importantly, this case provides a counterexample to a common belief according to which, say, an increase in the correlation of strategies due to information sharing is profitable if the reaction curves are upward-sloping.

4. For the cases mentioned under 3(a), I suggest a new explanation for the incentives to reveal private information which rests on two principles: (i) Letting rivals acquire better knowledge of their respective profit functions leads to a higher correlation of strategies, the profitability of which is determined by the slope of the reaction curves. (ii) Letting rivals acquire better knowledge of one's own profit function always increases the firm's own expected profits by inducing a change of the correlation of strategies in the direction which is profitable. The incentive to reveal information is then determined by the sum of these two effects.

2.2 General Model

In this section, a stochastic n-firm oligopoly model with private information is introduced at its most general level. In later sections, when I analyse particular aspects of information sharing, additional symmetry assumptions will have to be imposed.

We first discuss the main elements of the model: the State of Nature, private information, information sharing, and strategies and payoffs. Subsequently, explicit game formulations are given.

State of Nature: The State of Nature is denoted by the random variable $\boldsymbol{\tau} = (\tau_1, \ldots, \tau_n)'$, where τ_i is the deviation of either the marginal cost or the intercept of a linear demand function of firm i from its mean, depending on the type of uncertainty under consideration.[1] Note that for demand uncertainty, the intercepts may be different for each firm as well as for cost uncertainty. The variables τ_i are normal with zero mean, variance t_s and covariance $t_n \in [0, t_s]$. Let $\mathbf{I}$ denote the n-dimensional unit matrix, and define $\boldsymbol{\iota} = (1, 1, ..1)'$, and $\bar{\mathbf{I}} = \boldsymbol{\iota}\boldsymbol{\iota}' - \mathbf{I}$. Then the covariance matrix of $\boldsymbol{\tau}$ is given by $t_s\mathbf{I} + t_n\bar{\mathbf{I}} =: \mathbf{T}$.

[1] (i) The prime denotes transposition. (ii) For convenience, both the random variable and its realisations (hence particular States of Nature) are denoted by $\boldsymbol{\tau}$.

Private information: The State-of-Nature variable τ_i enters into firm i's profit function (see below), but is unknown to i. Instead, before setting a price or quantity, the firm – costlessly – receives a noisy signal y_i about τ_i as private information: $y_i := \tau_i + \eta_i$. The signal noise η_i is normal with zero mean, variance u_{ii}, and covariance $u_n \in [0, \min_i\{u_{ii}\}]$. Thus the covariance matrix of $\boldsymbol{\eta} = (\eta_1, \ldots, \eta_n)$ is $\text{diag}(u_{11}, \ldots, u_{nn}) + u_n\bar{\mathbf{I}} := \mathbf{U}$. Furthermore, $\boldsymbol{\tau}$ and $\boldsymbol{\eta}$ are assumed to be independent, which implies $\text{Cov}(\mathbf{y}) = \mathbf{T} + \mathbf{U} =: \mathbf{P}$.

The precision of firm i's signal is given by u_{ii}^{-1}: if $u_{ii} = 0$, firm i is perfectly informed about τ_i (e.g. its own cost); a positive u_{ii} implies a noisy signal, and for $u_{ii} = \infty$, y_i does not convey any information.

I follow Gal-Or (1985) in allowing that the signal errors η_i be correlated. For example, publicly accessible predictions about business cycles might enter into all y_i inducing a correlation not related to the true State of Nature. Hence the private signals may be correlated (i) due to a correlation of the components of the State of Nature and (ii) due to correlation of the signal errors.[2] [3] It is assumed throughout that the correlation of the signal errors does not exceed the correlation of the State-of-Nature components. This is stated more precisely in

Assumption COR: $t_n u_{ii} \geq t_s u_n \quad \forall i.$

Let $t_n/t_s =: \rho_\tau$ and $u_n/\sqrt{u_{ii}u_{jj}} =: \rho_\eta^{ij}$ (for $u_{ii}, u_{jj} > 0$) denote the correlation coefficients of $\boldsymbol{\tau}$ and $\boldsymbol{\eta}$, respectively. Then COR implies $\rho_\eta^{ij} \leq \rho_\tau$ for all i and j. If the u_{ii} are all equal, the two statements are equivalent. Assumption COR is automatically satisfied for all models in the literature. In the general model this assumption has to be made explicitly; its significance will become clear in the

[2] For analytical reasons it is required that the covariances between the signal errors are the same, which is a limitation of the model if the signal precisions are asymmetric. Thus we may either study the implications of correlated signal errors, assuming equal precisions, or analyse the effects of asymmetric precisions, assuming uncorrelated signal errors.

[3] In Gal-Or's (1985) model, however, the conditional correlation of the signal errors for a given State of Nature is *nonpositive*, an assumption for which Gal-Or does not provide an economic rationale.

next section.

Information revelation: Firms reveal their private information completely, partially, or not at all, to all other firms by means of a signal $\hat{y}_i := y_i + \xi_i$, where ξ_i is normal with zero mean and variance r_i. The ξ_i are independent of each other and of $\boldsymbol{\tau}$ and $\boldsymbol{\eta}$, hence for $\mathbf{r} = (r_1, \ldots, r_n)'$ and $\hat{\mathbf{y}} = (\hat{y}_1, \ldots, \hat{y}_n)'$ we have $\mathrm{Cov}(\boldsymbol{\xi}) = \mathrm{diag}(\mathbf{r})$ and $\mathrm{Cov}(\hat{\mathbf{y}}) = \mathbf{T} + \mathbf{U} + \mathrm{diag}(\mathbf{r}) =: \mathbf{Q}$. The variance r_i of the noise added to the true signal y_i expresses the revelation behaviour of firm i: for $r_i = 0$, y_i is completely revealed to the other firms; for $r_i = \infty$ a noisy signal with infinite variance is revealed, which is equivalent to concealing private information. For $0 < r_i < \infty$, private information is revealed partially: the signal y_i is distorted by the noise ξ_i, which reduces the informativeness of $\hat{y}_i$ according to the variance r_i. Note that y_i cannot be strategically distorted, since ξ_i and y_i are independent and ξ_i has zero mean. Hence apart from random noise, private information is (if at all) revealed truthfully, or equivalently, revealed information can be verified at no cost.[4]

Strategies and payoffs: Finally, we turn to the market structure of the model. Demand and cost functions are not explicit elements of the model. Rather, the profit functions are formulated directly. Each firm i controls the variable s_i, which is either the price of the good produced by i (Bertrand markets) or the quantity supplied (Cournot). The payoff for firm i is given by

$$\pi_i = \alpha_i(\tau_i) + \sum_{j \neq i} (\beta_n + \gamma_n \tau_i - \varepsilon s_i) s_j + (\beta_{ii} + \gamma_s \tau_i - \delta s_i) s_i, \tag{2.1}$$

where $\alpha_i(\tau_i)$ is any function of τ_i, and $\beta_{ii}, \beta_n, \gamma_s, \gamma_n, \delta,$ and ε are parameters. We assume that $\delta > 0$ and $\varepsilon \in (-\frac{1}{n-1}\delta, \delta]$.

The parametric profit function (2.1) suits a large range of standard oligopoly models, in particular, all types discussed in the information sharing literature (see Table 1 further below). This includes Cournot models with a linear demand system and linear or quadratic costs and Bertrand models with a linear demand

[4] This concept of partial revelation is due to Gal-Or (1985). A different, but qualitatively equivalent approach is used by Vives (1984) and Li (1985).

system and linear costs, in both cases for n firms producing heterogeneous goods. On the other hand, this generality also implies that there are no clear-cut economic interpretations of the parameters.

For all models with demand uncertainty (Cournot or Bertrand), γ_s equals 1, and for Cournot models with cost uncertainty, γ_s equals -1. In all these cases, γ_n equals zero. Hence (i) for Cournot competition, γ_s indicates the source of uncertainty, and (ii) only in the case of a Bertrand market with cost uncertainty, γ_n will take a nonzero value, the importance of which will be seen in later sections.

The linear-quadratic specification (2.1) arises from an underlying linear demand system with a coefficient matrix of the form $\mathbf{D} = \delta\mathbf{I} + \varepsilon\bar{\mathbf{I}}$ in both the Cournot and the Bertrand case. Such a demand system can be derived as the first-order condition of a representative consumer's maximisation of an appropriately defined utility function (cf. Vives 1984, Sakai and Yamato 1989), which in turn requires that the matrix $\mathbf{D}$ (or $\mathbf{D}^{-1}$, respectively) be positive definite, leading to the restriction on ε and δ stated above.

From (2.1), $\partial^2\pi_i/\partial s_i\partial s_j = -\varepsilon$. Hence for $\varepsilon > 0$ (e.g. Cournot with substitute goods or Bertrand with complements), we have a game of strategic substitutes, i.e. downward-sloping reaction curves; and strategic complements if $\varepsilon < 0$.

Game structures: We can now formulate the explicit game(s) that will be analysed. The model consists of the following stages:

(i) Firms decide on their revelation behaviour by setting r_i. I will consider two variants: (a) firms enter into a contract specifying that information shall be revealed completely, or not at all, i.e. $r_i = 0\ \forall i$ or $r_i = \infty\ \forall i$; (b) firms set the r_i's simultaneously, where I exclude partial revelation but allow for asymmetric behaviour, i.e. $r_i \in \{0, \infty\}\ \forall i$.

(ii) The State of Nature $\boldsymbol{\tau}$ is determined randomly. The players know the distribution of $\boldsymbol{\tau}$ but not its realisation.

(iii) Each firm i receives a private signal y_i. The distribution of $\mathbf{y}$ is common knowledge.

(iv) y_i is revealed completely, partially, or not at all, to all other firms by means of $\hat{y}_i$. The revelation behaviour is given by r_i, and $\mathbf{r}$ is known to all firms.

(v) Firms play the oligopoly game, i.e. each firm i sets the price/quantity s_i conditional on the information $\mathbf{z_i} := (y_i, \hat{\mathbf{y}}')'$ available to firm i.

Information structures: In Section 2.4, I will focus on some special cases of information structures to which the literature has restricted its attention. The first is the case of a Common Value, where according to the usual specification the State of Nature is a scalar entering into all firms' profits. Equivalently (since we are concerned with statistical decisions), we can assume that the n (identically distributed) components of the State of Nature are perfectly correlated, since then all τ_i are equal with probability one. We refer to this case as

Assumption CV (Common Value): $t_n = t_s =: t$.

All other cases, in which the State of Nature is a nondegenerate n-vector, have been referred to as "private-value" models. However, here we will distinguish two different kinds of those models. The first is the case where the components of the State of Nature are uncorrelated:

Assumption IV (Independent Values): $t_n = u_n = 0$,

where setting t_n to zero requires $u_n = 0$ (uncorrelated signal errors) because of COR. In fact, the work of Gal-Or (1986) is the only one in which assumption IV is made. In most of the other private-value models, any correlation between the State-of-Nature components is allowed for. But it is additionally assumed that firms receive signals without noise, i.e. acquire perfect knowledge about their "own" τ_i. We refer to this case as

Assumption PS (Perfect Signals): $u_{ii} = u_n = 0 \quad \forall i$.

Hence in this case, $\boldsymbol{\eta}$ degenerates to a zero distribution. Our separation of models classified as private-value models in the literature into two categories has two reasons: first, it seems more appropriate to refer to a "common value" and "independent values" as limit cases of the correlation of the τ_i lying between 0 and 1 rather than speak of a common value in the case of perfect correlation and of private values for any other case, including both independence and a correlation arbitrarily close to 1. Second and more importantly, only by taking the impact of signal noise into account, the apparent inconsistency pointed out in the Introduction between the results in common-value models and the results in certain "private-value" models can be explained, since the existence or nonexistence of signal noise is the only remaining difference between these types of models. The role of signal noise has not received any attention in previous work.

Almost all models in the literature are special cases of the model developed here, resulting by appropriately specifying the parameters.[5] These specifications are shown in Table 1. Note in particular that all models belong to one of the three classes CV, IV, and PS introduced above.

2.3 Nash Equilibrium of the Oligopoly Game

In this section, I derive the Bayesian Nash equilibrium of the oligopoly game. At this last stage, the revelation behaviour $\mathbf{r} = (r_1, \ldots, r_n)'$ is known to all firms, and each firm i has information $\mathbf{z_i} = (y_i, \hat{\mathbf{y}}')'$. The Bayesian Nash equilibrium $\mathbf{s}^*$ of this subgame is characterized by

$$s_i^*(\mathbf{z_i}) = \arg\max_{s_i \in IR^n} \mathrm{E}_{\tau,\eta_{-i}}[\pi_i(s_i, \mathbf{s}^*_{-\mathbf{i}} \mid \mathbf{z_i})] \quad (i = 1, \ldots, n),$$

[5] (i) Li (1985) and Shapiro (1986) have generalised the normality assumption by allowing for any distribution (e.g. one with compact support) for which all conditional expectations are affine functions of the given information variables. (ii) Kirby (1988) has studied information sharing agreements where nonrevealing firms are excluded from the pooled information. (iii) Hviid (1989) analyses information sharing between duopolists that are risk-averse. (iv) Shapiro (1986) considers (in our notation) τ_i's with different variances and the same correlation; and Sakai's (1986) perfect-signal duopoly model allows for arbitrary matrices $\mathbf{D}$ and $\mathbf{T}$. (v) The model of Novshek and Sonnenschein (1982) does not fit into our framework except for the uninteresting case of a common value *and* perfect signals (cf. the discussion in Clarke (1983)). These are the only exceptions.

Table 2.1: Previous models as special cases of the general model

model	n	$\alpha_i(\tau_i)$	ε	c_H	γ_n	β_{ii}	β_n	$\mathbf{T},\mathbf{U}$	$\mathbf{r}$
Clarke (1983)	n	0	δ	1	0	equal	0	CV, $\rho_\eta = 0$	$\mathbf{r} \in \{\mathbf{0}, \infty\boldsymbol{\iota}\}$
Fried (1984)	2	0	δ	1	0	diff.	0	PS	$r_i \in \{0, \infty\}$
Vives (1984)	2	0	any	1	0	equal	0	CV, $\rho_\eta = 0$	$r_i \in [0, \infty]$
Gal-Or (1985) (1)	n	0	δ	1	0	equal	0	CV, $u_n = -t_n$	$r_i \in [0, \infty]$
Gal-Or (1985) (2)	2	0	δ	1	0	equal	0	CV, $\rho_\eta \leq 0$	$r_i \in [0, \infty]$
Li (1985) (1)	n	0	δ	1	0	equal	0	CV, $\rho_\eta = 0$	$r_i \in [0, \infty]$
Li (1985) (2)	n	0	δ	-1	0	equal	0	PS	$r_i \in [0, \infty]$
Gal-Or (1986) (1)	2	0	any	-1	0	equal	0	IV	$r_i \in [0, \infty]$
Gal-Or (1986) (2)	2	$-\beta_{ii}\tau_i$	any	δ	ε	equal	0	IV	$r_i \in [0, \infty]$
Shapiro (1986)	n	0	δ	-1	0	equal	0	PS	$\mathbf{r} \in \{\mathbf{0}, \infty\boldsymbol{\iota}\}$
Sakai (1986)	2	0	any	1	0	diff.	0	PS	$r_i \in \{0, \infty\}$
Kirby (1988)	n	0	any	1	0	equal	0	CV, $\rho_\eta = 0$	$r_i \in \{0, \infty\}$
Sakai/ Yamato (1989)	n	0	any	-1	0	equal	0	PS	$\mathbf{r} \in \{0, \infty\boldsymbol{\iota}\}$

which leads to the reaction functions

$$s_i = \frac{1}{2\delta}\left[\beta_{ii} + \gamma_s E(\tau_i \mid \mathbf{z_i}) - \varepsilon \sum_{j \neq i} E(s_j \mid \mathbf{z_i})\right] \quad (i = 1, \ldots, n). \tag{2.2}$$

Here, expectations are formed over all random variables unknown at this stage, i.e. the State of Nature and the signal errors $\boldsymbol{\eta}_{-i}$ of the rival firms. Following the usual procedure, we derive the equilibrium strategies in two steps: first, we establish existence and uniqueness of an equilibrium with strategies s_i that are affine functions of $\mathbf{z_i}$. In the second step, the coefficients of these functions are computed.

Proposition 2.1 *There exists a unique Nash equilibrium of the oligopoly game for given information vectors* $\mathbf{z_i}$ $(i = 1, \ldots, n)$. *The equilibrium strategies* $s_i(\mathbf{z_i})$

are affine in $\mathbf{z_i}$*, i.e. for all* i*, there exist* a_i*,* $b_i \in \mathbb{R}$ *and* $\mathbf{c_i} \in \mathbb{R}^n$*, such that* $s_i = a_i + b_i y_i + \mathbf{c_i}\hat{\mathbf{y}}$.

For the proofs of all results in this chapter, see the Appendix (Section 2.6).

Having established linearity of the equilibrium strategies, we now compute the coefficients $a_i, b_i, \mathbf{c_i}$. To evaluate the first-order conditions (2.2), we first compute the conditional expectations $E(\tau_i \mid \mathbf{z_i})$ and $E(y_j \mid \mathbf{z_i})$.

Let $p_{ii} := t_s + u_{ii}\ \forall i$ and $p_n := t_n + u_n$ denote the variances and covariances of the signals y_i, respectively. Furthermore, define $m_i := (p_{ii} - p_n + r_i)^{-1}$ and $\mathbf{m} := (m_1, \ldots, m_n)'$. Finally, let $\mathbf{e_i}$ denote the i-th unit vector.

Proposition 2.2 *For given* $\mathbf{z_i}$*, the conditional expectations for* τ_i *and* y_j *are*

$$E(\tau_i \mid \mathbf{z_i}) = g_i y_i + \mathbf{\hat{g}_i'}\hat{\mathbf{y}} \quad \textit{and} \quad E(s_j \mid \mathbf{z_i}) = h_{ij} y_i + \mathbf{\hat{h}_{ij}'}\hat{\mathbf{y}} \quad (j \neq i), \quad \textit{where}$$

$$\begin{aligned} g_i &= \hat{t}/\hat{p}, & \mathbf{\hat{g}_i} &= (t_n p_{ii} - t_s p_n)(\mathbf{m} - m_i \mathbf{e_i})\hat{p}^{-1}, \\ h_{ij} &= p_n r_j m_j \hat{p}^{-1}, & \mathbf{\hat{h}_{ij}} &= (p_{jj} - p_n) m_j \mathbf{e_j} + h_{ij}(p_{ii} - p_n)(\mathbf{m} - m_i \mathbf{e_i}), \end{aligned}$$

$$\textit{and } \hat{t}_i = t_s + p_n(t_s - t_n)\sum_{j \neq i} m_j, \qquad \hat{p}_i = p_{ii} + p_n(p_{ii} - p_n)\sum_{j \neq i} m_j.$$

Setting $r_i = \infty$ for all i implies $\mathbf{m} = \mathbf{0}$ and $\mathbf{\hat{g}_i} = \mathbf{\hat{h}_{ij}} = \mathbf{0}$ for all i. Thus no use is made of the revealed signals $\hat{\mathbf{y}}$, which is equivalent to a situation without information sharing.[6]

The expression for $\mathbf{\hat{g}_i}$ makes the significance of assumption COR clear: since $t_n p_{ii} - t_s p_n = t_n u_{ii} - t_s u_n$, COR implies that the components of $\mathbf{\hat{g}_i}$ are nonnegative, which in turn ensures that a correlation of y_i and y_j is attributed to a correlation of τ_i and τ_j rather than to a correlation of the signal errors.

Substituting $E(s_j \mid \mathbf{z_i}) = a_j + b_j E(y_j \mid \mathbf{z_i}) + \mathbf{c_j'}\hat{\mathbf{y}}$ and the expressions from Proposition 2.2 in (2.2) yields

$$s_i(\mathbf{z_i}) = \frac{1}{2\delta}\left[\left(\beta_{ii} - \varepsilon\sum_{j \neq i} a_j\right) + \left(\gamma_s g_i - \varepsilon\sum_{j \neq i} b_j h_{ij}\right) y_i + \left(\gamma_s \mathbf{\hat{g}_i'} - \varepsilon\sum_{j \neq i} b_j \mathbf{\hat{h}_{ij}'} + \mathbf{c}_j'\right)\hat{\mathbf{y}}\right]. \tag{2.3}$$

[6] For infinite variances of r_i or u_{ii} I sometimes implicitly consider limit values of expressions of the kind found in Proposition 2.2. For example, "for $r_i = \infty$, $r_i m_i = 1$" is meant in the sense that $\lim_{r_i \to \infty} r_i m_i = 1$.

On the other hand, $s_i = a_i + b_i y_i + \mathbf{c}_i{}'\hat{\mathbf{y}}$. Identification of these coefficients with the corresponding terms in (2.3) leads to the main result of this section:

Proposition 2.3 *In the Bayesian Nash equilibrium of the oligopoly game each firm i $(i = 1, \ldots, n)$ has the strategy $s_i(\mathbf{z_i}) = a_i + b_i y_i + \mathbf{c}_i{}'\hat{\mathbf{y}}$, where*

$$\begin{aligned}
a_i &= \frac{1}{\tilde{d}}\left(\beta_{ii} - \frac{\varepsilon}{\hat{d}}\sum_{j=1}^{n}\beta_{ii}\right), \qquad b_i = \frac{\gamma_s}{v_i}\left(\hat{t}_i - \varepsilon p_n \frac{\sum_{j=1}^{n}\frac{r_i m_i \hat{t}_i}{v_i}}{1 + \varepsilon p_n \sum_{j=1}^{n}\frac{r_i m_i}{v_i}}\right), \\
\mathbf{c_i} &= \frac{2\delta}{\tilde{d}}\left[(p_{ii} - p_n)b_i - \frac{\varepsilon}{\hat{d}}\sum_{j=1}^{n}(p_{jj} - p_n)b_j \; - \; \frac{\gamma_s(t_s - t_n)}{\hat{d}}\right]\mathbf{m} \\
&\quad - \left[(p_{ii} - p_n)b_i - \frac{\gamma_s(t_s - t_n)}{\tilde{d}}\right] m_i \mathbf{e_i}, \text{ and} \\
\tilde{d} &= 2\delta - \varepsilon, \qquad \hat{d} = 2\delta + (n-1)\varepsilon, \qquad v_i = 2\delta\hat{p} - \varepsilon p_n r_i m_i.
\end{aligned}$$

The equilibrium strategies of the models of other works result as corollaries of Proposition 2.3: this applies for Clarke (1983), Fried (1984), Vives (1984) (Propositions 2, 2a); Gal-Or (1985, Theorems 1 and 2), Gal-Or (1986, Lemmas 1 and 2), Shapiro (1986), Li (1985, first model, Proposition 1), Kirby (1988), Sakai (1986), and Sakai and Yamato (1989).

An inspection of the expressions of Propositions 2.2 and 2.3 shows that although firm i does not use $\hat{y}_i$ for the expectations about τ_i or y_j, $\hat{y}_i$, its strategy s_i does depend on $\hat{y}_i$ since it enters into $E(s_j \mid \mathbf{z_i})$.

The strategies for the situation without information sharing follow as a limit case from Proposition 2.3 by setting $r_i = \infty$ for all i, which implies $\mathbf{c_i} = \mathbf{0}$ $\forall i$.

Using Proposition 2.3 we can derive the expected profits for the equilibrium of the oligopoly game (for simplicity denoted $E(\pi_i(\mathbf{s}))$), where expectations are formed for unknown $\mathbf{z_i}$ (i.e., before firms receive private information) but known revelation behaviour $\mathbf{r}$.

Proposition 2.4 *In the equilibrium given by Proposition 2.3, the expected profit for firm i is*

$$E(\pi_i(\mathbf{s})) = E(\alpha_i(\tau_i)) + \delta a_i^2 + \beta_n \sum_{j \neq i} a_j + \delta\, Var(s_i) + \gamma_n \sum_{j \neq i}(t_n b_j + \mathbf{c_j'}\mathbf{t_i}). \qquad (2.4)$$

For Cournot markets and for Bertrand markets with demand uncertainty, we have $\gamma_n = 0$, hence the last term in (2.4) vanishes. For most of the following sections we confine the analysis to these cases. Only Section 2.4.6 is devoted to the remaining case, Bertrand markets with cost uncertainty.

2.4 The incentives to share information

Several authors have noted that the incentives to share private information are largely determined by the change in the correlation of strategies induced by the pooling of information. However, it has never been treated analytically how this correlation is actually affected in different settings. Sections 2.4.1 and 2.4.2 address this question. We then study the two approaches to the determination of revelation behaviour introduced in Section 2.2: first, we analyse the incentives to completely pool information, compared with no pooling. Alternatively, we derive the equilibrium of the two-stage game where firms first independently decide on their revelation behaviour. A discussion in Section 2.4.5 draws the threads together. Finally, we turn to the case excluded for most of this chapter, Bertrand markets with cost uncertainty.

For the rest of the chapter we assume that $\beta_{ii} = \beta_s$ for all i. For most applications, this means that the firms have the same expected demand intercepts and marginal costs. Moreover, except for Section 2.4.6 we henceforth assume that $\gamma_n = 0$.

2.4.1 No-sharing case

As noted above, concealing of all private information corresponds to $r_i = \infty \ \forall i$, and from Proposition 2.3 we obtain (because of $m_i = 0$ and $r_i m_i = 1$)

$$b_i = \frac{\gamma_s t_s}{2\delta p_{ii} - \varepsilon p_n \sum\limits_{j \neq i} v_i / v_j} \quad \text{and} \quad \mathbf{c_i} = \mathbf{0} \ \forall i. \tag{2.5}$$

Without information sharing, therefore,

$$\mathrm{Var}(s_i) = p_{ii} b_i^2 \quad \text{and} \quad \mathrm{Cov}(s_i s_j) = b_i b_j \mathrm{E}(y_i y_j) = p_n b_i b_j \ \ (j \neq i). \tag{2.6}$$

Hence, for the correlation of these strategies, ρ_s^{ij}, we have $\rho_s^{ij} = p_n/\sqrt{p_{ii}p_{jj}}$, or if $p_{ii} := p_s \; \forall i$, $\rho_s = p_n/p_s$. Thus the correlation of s_i and s_j equals the correlation of the private signals y_i and y_j. Without sharing their private signals, players are not able to discriminate between the underlying State of Nature and the signal errors; therefore the correlation of the strategies does not depend on how the parameters of $\boldsymbol{\tau}$ and $\boldsymbol{\eta}$ enter into p_{ii} and p_n.

Next, we investigate how strategies and profits are influenced by the precision and correlation of the signals. From Proposition 2.4, changes in the information structure affect profits only inasmuch they affect $\mathrm{Var}(s_i)$. In the following, I will frequently use the notation $a \sim b$ to denote $\mathrm{sign}(a) = \mathrm{sign}(b)$.

Proposition 2.5 [7] *Without information sharing,*

$$(a)\frac{\partial b_i}{\partial p_{ii}} \sim -\gamma_s, \qquad (b)\frac{\partial E(\pi_i)}{\partial p_{ii}} < 0, \qquad (c)\frac{\partial b_i}{\partial p_{jj}} \sim \gamma_s \varepsilon, \qquad (d)\frac{\partial E(\pi_i)}{\partial p_{jj}} \sim \varepsilon.$$

Both the absolute value of b_i (the sign of which is determined by γ_s) and i's expected profit increase with the precision of y_i (parts a,b), whereas they are decreasing (increasing) in the precision of another firm's signal for strategic substitutes (complements) (parts c,d).

For the rest of the chapter, I assume that the private signals have equal precisions, i.e. $p_{ii} = p_s \; \forall i$. Then (2.5) implies $b_i = \gamma_s t_s \{p_s[2\delta + (n-1)\varepsilon\rho_y]\}^{-1} =: b$, where $\rho_y = p_n/p_s$ is the correlation of the signals. From $E(\pi_i(s)) \sim p_s b^2$ we immediately obtain (without proof)

Proposition 2.6 *In the completely symmetric model without information sharing,*

(a) $\partial E(\pi_i)/\partial p_s|_{\rho_y \, const.} < 0$, *i.e. for a given correlation of signals, a uniform increase in the precision of the signals increases expected profits;*

(b) $\partial E(\pi_i)/\partial \rho_y|_{p_s \, const.} \sim -\varepsilon$, *i.e for a given precision of signals, an increase in the correlation leads to higher expected profits for strategic complements and*

[7] The shorthand notation $E(\pi_i)$ refers to the expected profits for equilibrium strategies. Part (a) implies Lemma 1a in Vives (1984), and (b) implies Lemma 3a. From (b), Proposition 1 in Fried (1984) follows. Parts (c) and (d) imply parts of Lemmas 1b and 3b in Vives (1984).

to lower expected profits for strategic substitutes.

In contrast to the case of Proposition 2.5 (b), no relative information advantages of players are involved in result (a). Hence the precision of the private signal matters absolutely as well as in relation to the signals of the rival firms.

Some intuition on the well-known result (b) can be gained by considering a Cournot market with demand uncertainty (cf. Vives 1984): for a positive signal y_i, a higher correlation of signals implies a higher probability that the rival firms have received a high signal as well and supply a larger quantity. Since the reaction curves are downward-sloping, this induces a reduction of the own quantity s_i. As a result, i reacts less sensitively to y_i, which reduces the expected profit.

The result also explains parts (c) and (d) of Proposition 2.5: an exogenous increase in the precision of another player's signal (leaving the covariance unaffected) does not necessarily per se, i.e. because of an information advantage of the other firm, lead to a change of the expected profit, but rather through the increased correlation of strategies.[8] Hence the profitability depends on the sign of ε.

2.4.2 Complete pooling: correlation of strategies

Setting $r_i = 0$ for all i, we obtain the case of complete information sharing: all y_i are revealed without noise; all players have the same information. With $r_i m_i = 0$ and $m_i^{-1} = p_s - p_n$, Proposition 2.3 implies

$$b = \gamma_s \frac{t_s + (n-1)p_n \dfrac{t_s - t_n}{p_s - p_n}}{2\delta[p_s + (n-1)p_n]} \qquad \mathbf{c_i} = \left(b - \frac{\gamma_s}{\tilde{d}} \frac{t_s - t_n}{p_s - p_n}\right)\left(\frac{2\delta}{\hat{d}}\boldsymbol{\iota} - \mathbf{e_i}\right). \quad (2.7)$$

In the following, we usually focus on the cases CV, IV, and PS introduced above. For the case of a common value (CV), where $t_s = t_n = t$, the firms' strategies are identical and affine in the sample mean of the signals: $s_i = a + (2\delta b/\hat{d})\boldsymbol{\iota}'\mathbf{y}$. In

[8] Vives (1984) distinguishes the correlation effect and an information advantage of the rival firm, and argues that both affect expected profits negatively. However, it does not follow from his analysis that there exists a negative information advantage effect if the correlation of the signals is held constant.

the case of perfect signals (PS), where $p_s = t_s$ and $p_n = t_n$, the parameters of all random variables cancel out in (2.7), as all uncertainty has vanished (cf. Shapiro 1986).

Using (2.7) we can derive the variance and covariance of the equilibrium strategies, the sign of the correlation, and subsequently the direction of change with respect to the oligopoly without information pooling:

Proposition 2.7 *For CV, the correlation of equilibrium strategies always increases if information is completely pooled. For IV and PS, the correlation decreases for strategic substitutes and increases for strategic complements.*

(In particular, with independent values and strategic substitutes, information sharing leads to a *negative* correlation of previously uncorrelated strategies.)

Proposition 2.7 thus shows that the conjecture that the correlation increases with a common value and decreases with private values (cf. Li 1985, Gal-Or 1986) is correct for Cournot oligopolies with substitute goods, but not in general.

It should be noted that Proposition 2.7 does not require that $\gamma_n = 0$, i.e. it is valid for both Cournot and Bertrand, for demand and cost uncertainty.

2.4.3 Incentives to share I: contractual approach

The firms' incentives to enter into industry-wide contracts on information sharing are determined by difference of expected profits with and without information sharing. Using $E(\pi_i) \sim \text{Var}(s_i)$, the results of Section 2.4.1, and (2.18), we have

$$\begin{aligned} E(\pi_i^{CP}) - E(\pi_i^{NP}) = \delta\gamma_s^2 \Bigg\{ \frac{1}{\hat{d}^2} \bigg[\left(t_s + (n-1)p_n\tilde{t}\right) \frac{\bar{t}}{\bar{p}^n} \\ + \quad (n-1)\tilde{t} \left(\frac{4\delta^2 + (n-1)\varepsilon^2}{\tilde{d}^2}(t_s - t_n) - t_s \right) \bigg] - \frac{p_s t_s^2}{\hat{v}^2} \Bigg\}, \end{aligned} \tag{2.8}$$

where 'CP' and 'NP' refer to 'complete pooling' and 'no pooling', respectively; and $\bar{t} := t_s + (n-1)t_n$, $\bar{p}^k := p_s + (k-1)p_n$ and $\tilde{t} = (t_s - t_n)/(p_s - p_n)$. The sign of this difference does not depend on the sign of γ_s. Hence at least for Cournot models, the source of uncertainty – demand or cost – affects the signs of the strategies but is irrelevant for expected profits.

Instead of treating the IV and PS cases separately, we can derive more general results by taking an important similarity between these two cases into account: in both cases, firms do not acquire any new information about their τ_i's by the pooling of information: with perfect signals, firm i already knows τ_i, whereas with uncorrelated signals, it cannot infer anything about τ_i from the other firms' signals. In the model, this is reflected in the fact that in both cases, $\hat{\mathbf{g}}_\mathbf{i} = \mathbf{0}$. Evaluation of (2.8) leads to

Proposition 2.8 *If $\hat{\mathbf{g}}_\mathbf{i} = 0$, hence in particular for IV and PS, complete pooling is always profitable. For CV, pooling is profitable if and only if*

$$4\delta(\delta - \varepsilon)p_s - (n-1)\varepsilon^2(p_s + np_n) \quad > \quad 0$$

The expression on the l.h.s. is positive if $\mu := \varepsilon/\delta$ is less than $2/(n+1)$ and negative if μ is greater than $2(\sqrt{n}-1)/(n-1) < 1$, and otherwise depends on the magnitudes of p_s and p_n.

As corollaries follow the corresponding results of Clarke (1983), Fried (1984, Proposition 2), Li (1985, Proposition 2), Shapiro (1986, Theorem 1), Sakai (1986, Theorem 1), Kirby (1988, Proposition 2), and Vives (1984) (Proposition 5).

In a common-value Cournot oligopoly with sufficiently homogeneous goods (ε close to δ), complete sharing is unprofitable. In contrast, for small positive ε – corresponding to a large degree of product differentiation, or, for quadratic costs, to quickly increasing marginal costs (cf. Kirby 1988) – information sharing is profitable, as well as for negative ε (strategic complements).

The most important consequence of Proposition 2.8 is that with perfect signals, complete pooling is always profitable, regardless of any other parameters of the model. This result is in sharp contrast with the interpretations of Fried (1984), Shapiro (1986), Li (1985), Sakai (1986) and Sakai and Yamato (1989), who have attributed the profitability of information sharing to the "private-value" character of their models or the uncertainty about costs as opposed to demand uncertainty. Rather, the result is completely determined by the assumption that firms have perfect knowledge of their own costs, or in general, of their τ_i.

The proposition suggests that the unprofitability of information exchange in a homogeneous Cournot market with uncertainty about a common value is a rather exceptional case. Hence in general Clarke's (1983) argument that observing an agreement on information sharing may be taken as a prima-facie evidence of collusion does not apply.

2.4.4 Incentives to share II: noncooperative approach

We now analyse the two-stage game in which firms simultaneously decide on their revelation behaviour before playing the oligopoly game. While the "noncooperativeness" of this model structure obviously only relates to the revelation decisions, leaving their commitment character unaffected, studying the two-stage game can nevertheless yield important insights about the stability of information sharing arrangements. In particular, we will analyse under which circumstances firms have a dominant revelation strategy in the sense that they commit to a certain revelation behaviour (e.g. always to reveal the own signal) regardless of how the other firms decide, in anticipation of the equilibrium of the oligopoly game resulting from the first-stage decisions.

In this subsection, therefore, we allow for asymmetric revelation behaviour. However, we exclude partial revelation, i.e. each firm has to decide whether to reveal completely or not at all.[9]

Without loss of generality we assume that the first k players ($k \in \{0, \ldots, n\}$) reveal, whereas the last $n-k$ players conceal their information. A nonrevealing firm (for given k) has an incentive to reveal if

$$E(\pi_i(s_i^{R,k+1}, \mathbf{s}_{-\mathbf{i}}^{\mathbf{k+1}})) - E(\pi_i(s_i^{N,k}, \mathbf{s}_{-\mathbf{i}}^{\mathbf{k}})) > 0,$$

where $s_i^{R,k+1}$ denotes the strategy of a Revealing firm (increasing the number to $k+1$) and $s_i^{N,k}$ the strategy of a Nonrevealing firm (where the number of revealing firms remains k). If this inequality is valid for all k, i has a dominant strategy

[9] In analysing the two-stage game, Vives (1984) and Gal-Or (1985, 1986) allow for partial revelation, whereas this is excluded by Li (1985), Fried (1984), and Sakai (1986).

to reveal (in the sense explained above), and vice versa if the inequality is never fulfilled (cf. Li 1985).

Setting $r_i = 0$ for $i \in \{1, \ldots, k\}$ and $r_i = \infty$ for $i \in \{k+1, \ldots, n\}$ for given k, we can derive the equilibrium strategies for revealing and concealing firms from Proposition 2.3, calculate their variances, and compute the expected profits. This leads to one of the main results of this section:

Proposition 2.9 *For CV and $\hat{\mathbf{g}}_\mathbf{i} = 0$ (including IV and PS), there always exists a dominant revelation strategy (in the above sense). Information revelation is a strictly dominant strategy if $\hat{\mathbf{g}}_\mathbf{i} = 0$, and for CV and strategic complements. For CV and strategic substitutes, nonrevelation is a strictly dominant strategy.*

There are many corresponding results in the literature: Proposition 3 in Li (1985) and Proposition 3 in Fried (1984) follow as corollaries; and similar results are provided by Vives (1984), Gal-Or (1985, 1986), Li (1985), and Sakai (1986).

First of all, we observe that in all cases considered there are dominant revelation strategies. Furthermore, the result for PS complements Proposition 2.8: the results obtained by Fried (1984), Li (1985), and Sakai (1986) have little to do with cost uncertainty or "private values" but are determined by the mere assumption of perfect signals.

Comparing Propositions 2.8 and 2.9, we see that for most cases, the equilibrium of the two-stage game is efficient from the point of view of the firms. Only for CV with strategic substitutes and small ε (large degree of product differentiation) a Prisoner's Dilemma situation arises: complete sharing is profitable but does not occur in the two-stage game (cf. Vives 1984).

This, in turn, suggests that studying exclusionary disclosure rules (i.e. where only revealing firms have access to information revealed by others; such rules have been considered by Kirby (1988) and Shapiro (1986)) might not yield very interesting new insights, since "quid-pro-quo-agreements" (Kirby 1988) only become interesting in Prisoner's Dilemma situations where firms insist on the "quo". For exclusionary agreements among all n firms, of course, the results of Section 2.4.3 apply.

2.4.5 Discussion of the results

Excluding Bertrand markets with cost uncertainty from the analysis, we have shown: for CV and strategic complements, and for IV and PS in any case, complete information pooling is an efficient equilibrium of the two-stage game, regardless of all other parameters. For CV and strategic substitutes, no pooling is the equilibrium solution, which is efficient (inefficient) for a small (large) degee of product differentiation. Except for Gal-Or's (1986) Bertrand model with cost uncertainty, these statements summarise all results of the literature on the incentives to share information in symmetric models.

To explain the results of Sections 2.4.3 and 2.4.4, we start with the well-known common-value case (cf. Vives 1984). The pooling of information has two effects. First, each firm has better information about the prevailing market conditions; second, strategies are perfectly correlated. The first effect increases expected profits, whereas the profitability of the second effect depends on the slope of the reaction curves (cf. Proposition 2.5). For strategic complements, then, information sharing is unambiguously profitable. For strategic substitutes (say, a Cournot market with substitute goods), the correlation effect is negative. It outweighs the precision effect in the case of fairly homogeneous goods. With more differentiated goods, in contrast, the precision effect dominates (Proposition 2.8) since there is less intense competition, implying that the adverse effect of a higher correlation of strategies is smaller.[10]

In the noncooperative model, the decision to reveal only depends on the correlation effect, since the knowledge of the State of Nature is not influenced by the own revelation behaviour (cf. Proposition 2.2). This explains the difference between Propositions 2.8 and 2.9, which gives rise to a Prisoner's Dilemma.

While the distinction of a precision and a correlation effect is very useful in explaining the common-value case, it is of little use for the understanding of the IV and PS cases, as pointed out in the Introduction: recall that τ_i denotes

[10] For an alternative interpretation in the model with quadratic costs, see Kirby (1988).

the component of the State of Nature which enters into firm i's profit function. First, information sharing cannot improve firm i's information on τ_i (cf. 4.3), and it is not clear why it would benefit from improved information about the other τ_j *as such.* In fact, Fried (1984) and Sakai (1985) provide examples in which firms prefer never to receive any signals about the rival's profit function. Second, while the change in the correlation of strategies is clearly important, it is endogenous and not easily predicted. In particular, our results show that in the IV and PS cases information sharing changes the correlation of strategies exactly in the direction that is profitable for the firms! This, of course, undermines the explanatory power of the correlation effect.

Therefore, we proceed to explain our results in terms of two rather different, and more general, effects in the change of strategies due to information sharing: "direct adjustments", due to an improved knowledge of the own τ_i, and "strategic adjustments", due to an improved knowledge of the rival firms' information and hence their actions.[11]

These effects can be readily identified analytically: while the direct adjustment is directly related to the magnitude of the parameters of $\hat{\mathbf{g}}_\mathbf{i}$, the strategic adjustments are determined by the parameters of $\hat{\mathbf{h}}_{\mathbf{ij}}$ (cf. (2.2) and Proposition 2.2).

The significance of our new distinction is immediately clear: in both cases IV and PS there are only strategic adjustments, since $\hat{\mathbf{g}}_\mathbf{i} = \mathbf{0}$. From Proposition 2.9 we may thus conclude that for IV and PS, unilateral revelation of information to the other firms is profitable because and as long as this induces only strategic adjustments by the rival firms.

Our distinction also sheds new light on the common-value case: while strategic adjustments always alter the correlation of strategies in the direction profitable for the firms (Propositions 2.6b, 2.7, and 2.8), direct adjustments always lead to a higher correlation. Thus with strategic complements, both adjustments are

[11] This terminology is borrowed from Fried (1984), who uses the terms "direct adjustments" and "counteradjustments" in the same way, but in a slightly different context.

profitable, whereas with strategic substitutes, the negative effect of highly correlated strategies may prevail. For CV, in particular, the components of $\hat{\mathbf{g}}_i$ have their maximal value (cf. Proposition 2.3), implying maximal direct adjustments.

Turning to intermediate cases between IV and CV, in markets with strategic substitutes firms face a trade-off: a firm has an incentive to reveal its private information, as long as this does not significantly improve other firms' knowledge of their τ_j which would induce direct adjustments by these firms and thereby lead to more intense competition (cf. Fried 1984, Proposition 4]).

2.4.6 Bertrand markets with cost uncertainty

We briefly turn to Bertrand markets with cost uncertainty. Consider the simplest example with the demand function $q_i = a - \delta p_i - \varepsilon \sum_{j \neq i} p_j$ and a random marginal cost c_i, for simplicity with zero mean. In terms of the profit function (2.1), $\gamma_n = \varepsilon$, which is the coefficient for the product $c_i \sum_{j \neq i} p_j$. Such terms vanish in each of the other cases I have been considering, but play an important role in this case. Therefore, Bertrand competition with cost uncertainty is structurally different from the other three cases.

With $\gamma_n \neq 0$, the last term in (2.4) does not vanish. As a consequence, the analysis becomes considerably more complicated, and the results are much more ambiguous, than in the other cases. Therefore, I will only summarise the results without presenting the formal analysis in full detail. As to the method of analysis, it suffices to analyse the last term in (2.4) under different informational settings, proceeding exactly as in the previous sections, and then combine the results with the corresponding results derived in 4.1-4.4.

In contrast to the simple results in the previous sections, the profitability of industry-wide contracts on information sharing in general depends on the magnitudes of δ, ε, and n. Similarly, the profitability of unilateral information revelation depends on the these parameters. Moreover, in general there do not even exist dominant revelation strategies.

Only in the case of independent values does a dominant revelation strategy exist. This strategy depends on the difference between expected profits for unilateral revelation vs. concealing, which has the same sign as $-4(2\delta-\varepsilon)-(n-1)\varepsilon(4\delta-3\varepsilon)$. If $\varepsilon > 0$ or $n = 2$, to conceal information is a dominant revelation strategy. This was shown by Gal-Or (1986) for the duopoly case. However, for negative ε and $n \to \infty$, revealing becomes a dominant strategy. Thus, even when a dominant revelation strategy exists, whether this strategy involves revelation or not depends on the specific parameters. Results obtained for duopolies do not extend to larger markets.

To see why this case is so different, contrast the profit functions for a Bertrand duopoly for demand uncertainty, $\pi_i = p_i(a-\tau_i-\delta p_i-\varepsilon p_j)$ and for cost uncertainty, $\pi_i = (p_i-\tau_i)(a-\delta p_i-\varepsilon p_j)$. For substitutes, ε is negative. In both cases, a positive τ_i will affect the profit negatively.

Now fix p_i as a random variable and consider how the expected profits in both cases depend on the rival's strategy p_j. For demand uncertainty, the term $-\varepsilon \mathrm{E}(p_i p_j)$ enters into the expected profits. Hence, we obtain the well-known result that expected profits are increasing in the correlation of the firms' strategies. With cost uncertainty, in contrast, we get $-\varepsilon \mathrm{E}(p_i p_j)+\varepsilon \mathrm{E}(\tau_i p_j)$. With information sharing, the second term counterbalances the first, since (with $\gamma_s = \delta > 0$) p_i and τ_i are positively correlated. Considering the effect of this second term, therefore, it is not surprising that the profitability of information sharing depends on the parameters of the specific model.

As noted above, Proposition 2.7 also covers Bertrand markets with cost uncertainty. The important conclusion is that while information sharing leads to the unambiguous change in the correlation of strategies stated in Proposition 2.7 (depending on the information structure), only for the three cases considered in Sections 2.4.4 and 2.4.5 it is true that an increase in the correlation of the firms' strategies is profitable for strategic complements, contrary to what is usually believed. In contrast to Gal-Or's (1986) interpretation of her result, therefore, the nonprofitability of information sharing in the Bertrand duopoly with cost uncertainty does not arise *because of a decrease* in the correlation of strategies, but rather *despite an increase.*

What is remarkable, then, is not the ambiguity observed here, but the simplicity of the results of the previous sections. This simplicity hinges on a simple relationship between expected profits and the variances and covariances of the equilibrium strategies which does not exist in the case considered here.

2.5 Concluding Remarks

Previous work has fostered the impression that the incentives to share information depend delicately on the details of the model. In contrast, I have shown - building on more general results - that the results for the majority of the specific models can be summarised in a very simple way.

Our analysis suggests that some generalising interpretations of those results found in other works are invalid: (i) As I have argued in 4.5, the assertion that one major determinant encouraging firms to exchange information is an improvement of the information about market conditions, is valid only as far as information about *own* demand or cost is concerned, but then does not apply to models in which firms cannot improve this information, viz. in independent-value or perfect-signal models (which, in fact, comprise at least one half of those considered in the literature). (ii) The assertion that the other major determinant of the profitability of information sharing is the induced change in the correlation of strategies is unhelpful, since this change in the correlation is itself endogenous and not easily predicted without explicit formal analysis. In particular, we have shown that for independent-value or perfect-signal models, the correlation always changes in the direction which is profitable for the firms, although this direction depends on the details of the profit function.

Our new alternative interpretation, which applies to all information structures and to all market types considered except for Bertrand markets with cost uncertainty, rests on two separate effects which determine the incentives to reveal information: 1. Letting rivals acquire better knowledge of their respective profit functions leads to a higher correlation of strategies, the profitability of which is determined by the slope of the reaction curves. 2. Letting rivals acquire better knowledge of one's own profit function is *always* profitable.

An analysis of the welfare effects of information sharing, not included here, leads to less clear-cut results than with the incentives for firms to share information. In many cases, the direction of change of consumer surplus and total welfare depends on the magnitudes of the parameters of the model.

The intuitive conjecture, often found in the nonformal literature, that without collusion information sharing is socially beneficial, can by and large be supported, as far is overall welfare is concerned. However, it fails to take into account the impact of a change in the correlation of strategies on profits and the effect on consumer surplus. In general, producers and consumers have conflicting interests, making a weighting of these interests necessary (cf. Shapiro 1986).

The understanding of the role of information in oligopoly could be further improved by studying asymmetric information structures in more detail. Situations where firms have differently precise private information have been analysed by Clarke (1983), Fried (1984), and Sakai (1986).

While in this chapter the incentives for firms to *reveal* private information have been emphasized, two other related issues are the incentives to *acquire* (costly) information about the own profit function, and finally to *receive* information about other firms. The first has been pursued by Li, McKelvey, and Page (1987), Vives (1988) and recently by Hwang (1993), the second by Fried (1984), Sakai (1985), and Jin (1992).

Despite the generality of the structure used, the present model might be considered restrictive in some respects. It rests on linear-quadratic profit functions, normally distributed random variables, and various symmetry assumptions. Commitment to a revelation strategy and truthtelling are imposed by assumption. In this respect this chapter is no more general than previous work. This is certainly a justified criticism which calls for the development of still more general frameworks which allow for an assessment of the robustness of our results.

But while recent research has moved on to generalise certain elements of the earlier models, or - probably more fruitfully - to develop strategically more sophisticated models which build on the methodological criticisms raised against the standard models (Zvi (1993) makes a step in this direction), all the old questions which gave rise to this line of research have been left unanswered, as forcefully argued by Vives (1990). This chapter is a contribution to fill this gap; i.e. though adhering to a restrictive framework, it does lead to a general theory for a large class of models which has been the focus of research for ten years, thus providing a benchmark for departures from this framework.

2.6 Appendix: Proofs

Most proofs involve very tedious but straightforward algebra which has been omitted here. The details can be found in Raith (1983).

Proof of Proposition 2.1: Equations (2.2) can be obtained as the first-order conditions of an appropriately defined team-decision problem in the sense of Radner (1962). Similarly, all higher order conditions for the Nash equilibrium and the solution of the team decision problem are identical. Therefore, we immediately obtain the result by applying Theorem 5 of Radner (1962), where in particular the assumption of normal distributions and the positive definiteness of $\mathbf{D}$ ensure that all assumptions of that theorem are satisfied.[12] ■

Proof of Proposition 2.2: 1. The proofs of the results in Section 2.3 make use of the following result: For matrices $\mathbf{A}$ ($p \times p$, nonsingular), $\mathbf{U}, \mathbf{V}$ ($q \times p$), and $\mathbf{S}$ ($q \times q$) we have

$$(\mathbf{A} + \mathbf{U}'\mathbf{S}\mathbf{V})^{-1} = \mathbf{A}^{-1} - \mathbf{A}^{-1}\mathbf{U}'\mathbf{S}(\mathbf{S} + \mathbf{S}\mathbf{V}\mathbf{A}^{-1}\mathbf{U}'\mathbf{S})^{-1}\mathbf{S}\mathbf{V}\mathbf{A}^{-1} \qquad (2.9)$$

(see Madansky (1976, p. 9]). Now let $a, b \in \mathbb{R}, \mathbf{a} \in \mathbb{R}^n$, and $\bar{\mathbf{a}} = \left(\frac{1}{a_1}, \ldots, \frac{1}{a_n}\right)'$. Then from (2.9) it follows that

$$[\mathrm{diag}(\mathbf{a}) + b\boldsymbol{\iota}\boldsymbol{\iota}']^{-1} = \mathrm{diag}(\bar{\mathbf{a}}) - \frac{b}{1 + b\boldsymbol{\iota}'\bar{\mathbf{a}}}\bar{\mathbf{a}}\bar{\mathbf{a}}' \quad \text{and} \qquad (2.10)$$

$$[a\mathbf{I} + b\boldsymbol{\iota}\boldsymbol{\iota}']^{-1} = a^{-1}\mathbf{I} - (a(a+bn))^{-1}b\boldsymbol{\iota}\boldsymbol{\iota}'. \qquad (2.11)$$

2. Since $\boldsymbol{\tau}$, $\boldsymbol{\eta}$, and $\boldsymbol{\xi}$ are independent, we have $\mathrm{Var}(\boldsymbol{\tau}) = \mathrm{Cov}(\boldsymbol{\tau}, \mathbf{y}) = \mathrm{Cov}(\boldsymbol{\tau}, \hat{\mathbf{y}}) = \mathbf{T}$, $\mathrm{Var}(\mathbf{y}) = \mathrm{Cov}(\mathbf{y}, \hat{\mathbf{y}}) = \mathbf{P}$ and $\mathrm{Var}(\hat{\mathbf{y}}) = \mathbf{Q}$. Writing $\mathbf{T}$ and $\mathbf{P}$ as rows of column vectors, $\mathbf{T} = (\mathbf{t_1}, \ldots, \mathbf{t_n})$ and $\mathbf{P} = (\mathbf{p_1}, \ldots, \mathbf{p_n})$, we obtain

$$\mathrm{E}(\tau_i \mid (y_i, \hat{\mathbf{y}}')') = (t_s, \mathbf{t_i}')\mathbf{V}(y_i, \hat{\mathbf{y}}')' \quad \text{and} \quad \mathrm{E}(y_j \mid (y_i, \hat{\mathbf{y}}')') = (p_n, \mathbf{p_j}')\mathbf{V}(y_i, \hat{\mathbf{y}}')',$$

where

$$\mathbf{V} = \begin{pmatrix} p_{ii} & \mathbf{p_i}' \\ \mathbf{p_i} & \mathbf{Q} \end{pmatrix}^{-1} = \begin{pmatrix} C_1 & -C_1\mathbf{p_i}'\mathbf{Q}^{-1} \\ -C_1\mathbf{Q}^{-1}\mathbf{p_i} & \mathbf{Q}^{-1} + C_1\mathbf{Q}^{-1}\mathbf{p_i}\mathbf{p_i}'\mathbf{Q}^{-1} \end{pmatrix}$$

[12] See Basar and Ho (1974) for a similar application of Radner's theorem to a duopoly model and Vives (1988) for the application to both oligopolies and competitive markets.

and $C_1 = (p_{ii} - \mathbf{p_i'Q^{-1}p_i})^{-1}$ (cf. Theil (1971), p.17-18]). From this we obtain

$$E(\tau_i \mid (y_i, \hat{\mathbf{y}}')') = g_i y_i + \hat{\mathbf{g}}_i'\hat{\mathbf{y}} \quad \text{and} \quad E(y_j \mid (y_i, \hat{\mathbf{y}}')') = h_{ij} y_i + \hat{\mathbf{h}}_{ij}'\hat{\mathbf{y}}, \text{where}$$

$$\begin{aligned} g_i &= \frac{t_s - \mathbf{t_i'Q^{-1}p_i}}{p_{ii} - \mathbf{p_i'Q^{-1}p_i}}, & \hat{\mathbf{g}}_i &= \mathbf{Q}^{-1}(\mathbf{t_i} - g_i\mathbf{p_i}), \\ h_{ij} &= \frac{p_n - \mathbf{p_j'Q^{-1}p_i}}{p_{ii} - \mathbf{p_i'Q^{-1}p_i}}, & \hat{\mathbf{h}}_{ij} &= \mathbf{Q}^{-1}(\mathbf{p_j} - h_{ij}\mathbf{p_i}). \end{aligned}$$

With $\mathbf{m}$ as defined in the main text it follows that $\mathbf{Q} = \text{diag}(m_1^{-1}, \dots, m_n^{-1}) + p_n\boldsymbol{\iota\iota}'$, and hence from (2.10): $\mathbf{Q^{-1}} = \text{diag}(\mathbf{m}) - (1 + p_n \sum_{j=1}^n m_j)^{-1} p_n \mathbf{mm}'$. Writing $\mathbf{t_i} = t_n\boldsymbol{\iota} + (t_s - t_n)\mathbf{e_i}$ etc., straightforward calculations then lead to the expressions stated in the proposition. ∎

Proof of Proposition 2.3: 1. From (2.3), $2\delta a_i = \beta_{ii} - \varepsilon \sum_{j \neq i} a_j$ or $(2\delta - \varepsilon)a_i = \beta_{ii} - \varepsilon \sum_{j=1}^n a_j \ \forall i$, from which the expression for α_i follows in an obvious way.

2. From (2.3), $2\delta b_i = \gamma_s g_i - \varepsilon \sum_{j \neq i} b_j h_{ij} \ \forall i$, which can be written as

$$\begin{pmatrix} 2\delta & \varepsilon h_{12} & \cdots & \varepsilon h_{1n} \\ \varepsilon h_{21} & 2\delta & & \vdots \\ \vdots & & \ddots & \vdots \\ \varepsilon h_{n1} & \cdots & \cdots & 2\delta \end{pmatrix} \begin{pmatrix} b_1 \\ \vdots \\ b_n \end{pmatrix} = \gamma_s \begin{pmatrix} g_1 \\ \vdots \\ g_n \end{pmatrix}. \tag{2.12}$$

Define $\boldsymbol{\theta} = (\hat{t}_1, \dots, \hat{t}_n)$ and $\tilde{\mathbf{m}}' = (r_1 m_1, \dots, r_n m_n)$. Since $g_i = \hat{t}_i / \hat{p}_i$ and $h_{ij} = p_n r_j m_j / \hat{p}_i$, it follows that

$$[\text{diag}(\mathbf{v}) + \varepsilon p_n \boldsymbol{\iota}\tilde{\mathbf{m}}']\,\mathbf{b} = \gamma_s \boldsymbol{\theta} \quad \text{or} \quad \mathbf{b} = \gamma_s\,[\text{diag}(\mathbf{v}) + \varepsilon p_n \boldsymbol{\iota}\tilde{\mathbf{m}}']^{-1}\,\boldsymbol{\theta}.$$

Using (2.9) we obtain

$$[\text{diag}(\mathbf{v}) + \varepsilon p_n \boldsymbol{\iota}\tilde{\mathbf{m}}']^{-1} = \left(\mathbf{I} - \frac{\varepsilon p_n}{(1 + \varepsilon p_n \tilde{\mathbf{m}}'\bar{\mathbf{v}})}\bar{\mathbf{v}}\tilde{\mathbf{m}}'\right) \text{diag}(\bar{\mathbf{v}}),$$

where $\bar{\mathbf{v}} := (v_1^{-1}, \dots, v_n^{-1})$. The i-th row of $\mathbf{I} - \varepsilon p_n \bar{\mathbf{v}}\tilde{\mathbf{m}}' / (1 + \varepsilon p_n \tilde{\mathbf{m}}'\bar{\mathbf{v}})$ is $\mathbf{e_i}' - (\varepsilon p_n / \chi)\bar{v}_i \tilde{\mathbf{m}}'$, which leads to the expression stated in the proposition.

4. From (2.3),

$$2\delta\mathbf{c_i} + \varepsilon \sum_{j \neq i} \mathbf{c_j} = \gamma_s \hat{\mathbf{g}}_i - \varepsilon \sum_{j \neq i} b_j \hat{\mathbf{h}}_{ij} \quad \forall i. \tag{2.13}$$

Now define $\mathbf{C} := (\mathbf{c_1}, \ldots, \mathbf{c_n})$ and

$$\mathbf{Z} := \left(\gamma_s \hat{\mathbf{g}}_1 - \varepsilon \sum_{j \neq 1} b_j \hat{\mathbf{h}}_{1j}, \ldots, \gamma_s \hat{\mathbf{g}}_n - \varepsilon \sum_{j \neq n} b_j \hat{\mathbf{h}}_{nj}\right).$$

Then the system (2.13) can be written as $\mathbf{C}(\tilde{d}\mathbf{I} + \varepsilon\bar{\mathbf{I}}) = \mathbf{Z}$. Using (2.11) we obtain

$$\mathbf{c_i} = \mathbf{Z}\frac{1}{\tilde{d}}\left(\mathbf{e_i} - \frac{\varepsilon}{\hat{d}}\boldsymbol{\iota}\right) = \frac{1}{\tilde{d}}\left[\gamma_s \hat{\mathbf{g}}_i - \varepsilon \sum_{j \neq i} b_j \hat{\mathbf{h}}_{ij} - \frac{\varepsilon}{\hat{d}} \sum_{j=1}^{n} \left(\gamma_s \hat{\mathbf{g}}_j - \varepsilon \sum_{k \neq j} b_k \hat{\mathbf{h}}_{jk}\right)\right]. \tag{2.14}$$

Substituting the expressions for $\hat{\mathbf{g}}_i$ and $\hat{\mathbf{h}}_{ij}$ given in Proposition 2.2 eventually yields the result given in the proposition. ■

Proof of Proposition 2.4: According to (2.1), firm i's profit is

$$\pi_i = \alpha_i(\tau_i) + (\beta_n + \gamma_n \tau_i) \sum_{j \neq i} s_j + \left[\beta_{ii} + \gamma_s \tau_i - \varepsilon \sum_{j \neq i} s_j - \delta s_i\right] s_i. \tag{2.15}$$

Since i knows $\mathbf{z_i}$ when she determines s_i, $\mathrm{E}(\pi_i(\mathbf{s}))$ equals

$$\mathrm{E}(\alpha_i(\tau_i)) + E\left[(\beta_n + \gamma_n \tau_i) \sum_{j \neq i} s_j\right] + E_{y_i, \hat{y}}\left[E\left(\beta_{ii} + \gamma_s \tau_i - \varepsilon \sum_{j \neq i} s_j - \delta s_i \mid \mathbf{z_i}\right) s_i\right]. \tag{2.16}$$

Because of (2.2), the last term in (2.16) reduces to $\delta E(s_i^2)$, since s_i is an equilibrium strategy; and then $\mathrm{E}(s_i^2) = E^2(s_i) + \mathrm{Var}(s_i) = a_i^2 + \mathrm{Var}(s_i)$. For the second term in (2.16) we get

$$E\left[(\beta_n + \gamma_n \tau_i) \sum_{j \neq i} s_j\right] = \sum_{j \neq i} (\beta_n a_j + \gamma_n b_j t_n + \gamma_n \mathbf{c}_j' \mathbf{t_i}),$$

where $\mathbf{t_i}$ is the i-th column vector of $\mathbf{T}$. Then the expression stated in the proposition results immediately.

■

Proof of Proposition 2.5: First, the restriction $\varepsilon > -(n-1)^{-1}\delta$ implies that $1 + \varepsilon p_n \sum_{j=1}^{n} (1/v_j) > 0$ and hence $b_i \sim \gamma_s$. From (2.5), we have

$$\frac{\partial b_i}{\partial p_{ii}} = -\frac{2\delta b_i}{v_i}\left(1 - \frac{\varepsilon p_n / v_i}{1 + \varepsilon p_n \sum_{j=1}^{n} (1/v_j)}\right).$$

The two previous results then imply part (a). Part (b) follows from $\partial E(\pi(s))/\partial p_{ii} \sim \partial p_{ii} b_i^2 / \partial p_{ii}$, $\partial(p_{ii} b_i)/\partial p_{ii} \sim -\gamma_s$, $p_{ii} b_i \sim b_i \sim \gamma_s$, and part (a). Parts (c) and (d) can then be derived in a straightforward way. ■

Proof of Proposition 2.7: 1. From $E(y_i^2) = p_s, E[(\mathbf{c}_i'\hat{\mathbf{y}})^2] = E(\mathbf{c}_i'\hat{\mathbf{y}}\hat{\mathbf{y}}'\mathbf{c}_i) = \mathbf{c}_i'\mathbf{Q}\mathbf{c}_i$ and $E(by_i\hat{\mathbf{y}}'\mathbf{c}_i) = b\mathbf{p}_i'\mathbf{c}_i$ we obtain

$$\mathrm{Var}(s_i^{CP}) = E[(by_i + \mathbf{c}_i'\hat{\mathbf{y}})^2] = p_s b^2 + \mathbf{c}_i'\mathbf{Q}\mathbf{c}_i + 2b\mathbf{p}_i'\mathbf{c}_i. \tag{2.17}$$

Using (2.7) and $\mathbf{Q} = (p_s - p_n)\mathbf{I} + p_n\boldsymbol{\iota}\boldsymbol{\iota}'$ and (2.7), (2.17) leads to

$$\mathrm{Var}(s_i^{CP}) = \frac{c_H^2}{\hat{d}^2}\left[t_s\frac{\bar{t}}{\bar{p}^n} + (n-1)\tilde{t}\left(p_n\frac{\bar{t}}{\bar{p}^n} - t_s + \frac{4\delta^2 + (n-1)\varepsilon^2}{\tilde{d}^2}(t_s - t_n)\right)\right]. \tag{2.18}$$

where $\bar{t} := t_s + (n-1)t_n$, $\bar{p}^k := p_s + (k-1)p_n$ and $\tilde{t} = (t_s - t_n)/(p_s - p_n)$. Proceeding in the same way, we obtain a similar expression for the covariance:

$$\mathrm{Cov}(s_i^{CP}, s_j^{CP}) = \frac{c_H^2}{\hat{d}^2}\left[t_s\frac{\bar{t}}{\bar{p}^n} + (n-1)\tilde{t}\left(p_n\frac{\bar{t}}{\bar{p}^n} - t_n\right) - \frac{4\delta^2 + (n-1)\varepsilon^2}{\tilde{d}^2}(t_s - t_n)\tilde{t}\right]. \tag{2.19}$$

2. Let $\rho_s = \mathrm{Cov}(s_i, s_j)/\mathrm{Var}(s_i)$ denote the correlation of equilibrium strategies, where ρ_s^{CP} and ρ_s^{NP} refer to the complete-pooling and no-pooling cases, respectively.

CV: Trivial since for complete sharing the strategies are identical.

IV: The result follows immediately from $\mathrm{Cov}(s_i^{CP}, s_j^{CP}) \sim \varepsilon$ and $\rho_s^{NP} = 0$.

PS: With $\rho_s^{CP} - \rho_s^{NP} \sim p_s\mathrm{Cov}(s_i, s_j) - p_n\mathrm{Var}(s_i)$, which using (2.18) and (2.19) can be shown to be $\sim -\varepsilon$. ■

Proof of Proposition 2.8: 1. CV: From (2.8) we obtain

$$E(\pi_i^{CP}) - E(\pi_i^{NP}) \sim 4\delta(\delta - \varepsilon) - (n-1)\varepsilon^2 - n(n-1)\varepsilon^2\frac{p_n}{p_s}.$$

Since $p_n/p_s \in [0,1]$, the last expression is necessarily positive if $4\delta(\delta-\varepsilon) - (n-1)\varepsilon^2 > n(n-1)\varepsilon^2$ and negative if $4\delta(\delta-\varepsilon) - (n-1)\varepsilon^2 < 0$. Calculating the critical values of ε for these inequalities then leads to the result stated in the proposition.

2. $\hat{\mathbf{g}}_i = \mathbf{0}$: Since $t_n p_s = t_s p_n$, we obtain from (2.8):

$$E(\pi_i^{CP}) - E(\pi_i^{NP}) = (n-1)\delta\varepsilon^2\frac{c_H^2}{\tilde{d}^2\hat{d}^2\hat{v}^2}(t_s - t_n)\bar{p}^n\tilde{t}\left[\hat{v}(\tilde{d} + \hat{d}) + \tilde{d}\hat{d}p_s\right] > 0.$$

■

Proof of Proposition 2.9: 1. We first calculate for a given k the coefficients b and $\mathbf{c}_i$ for revealing and nonrevealing firms in the general case. Using R as an index for a revealing firm i ($i = 1, \ldots, k$) and N as an index for a nonrevealing firm i ($i = k+1, \ldots, n$), we get from Proposition 2.3

$$
\begin{aligned}
b^N &= c_H \frac{t_s + kp_n\tilde{t}}{2\delta(p_s + kp_n) + (n-k-1)\varepsilon p_n}, && (2.20)\\
b^R &= \frac{c_H}{2\delta \bar{p}^k}\left[\frac{2\delta(p_s + kp_n) - \varepsilon p_n}{\tilde{v}_k}(t_s + kp_n\tilde{t}) - p_n\tilde{t}\right] \quad \text{and} && (2.21)\\
\mathbf{c_i} &= \frac{2\delta}{\tilde{d}}\left(b_i - \frac{\varepsilon}{\hat{d}}\sum_{j=1}^{n} b_j - \frac{c_H}{\hat{d}}\tilde{t}\right)\boldsymbol{\iota}_k - \left(b_i - \frac{c_H}{\tilde{d}}\tilde{t}\right)\mathbf{e_i^R}, && (2.22)
\end{aligned}
$$

where $\tilde{v}_k := 2\delta(p_s + kp_n) + (n-k-1)\varepsilon p_n$, $\boldsymbol{\iota}_k$ is a vector with ones in the first k components and zeros in the last $n-k$, and $\mathbf{e_i^R} = \mathbf{e_i}$ if $i \le k$, and $\mathbf{e_i^R} = \mathbf{0}$ otherwise. 2. CV: Equations (2.20) to (2.22) imply

$$
s_i^R = a + \frac{c_H t_s}{\hat{d}\bar{p}^k}\,\boldsymbol{\iota}_k'\mathbf{y_k} \quad \text{and} \quad s_i^N = a + \frac{c_H t_s}{\tilde{v}}y_i + \frac{c_H t_s}{\hat{d}\bar{p}^k}\frac{2\delta\bar{p}^k - k\varepsilon p_n}{\tilde{v}}\boldsymbol{\iota}_k'\mathbf{y_k},
$$

where $\mathbf{y_k}$ is the vector of signals $\mathbf{y}$ with zeros in the last $(n-k)$ components. Noting that $E[(\boldsymbol{\iota}_k'\boldsymbol{y})^2] = \boldsymbol{\iota}_k'\boldsymbol{P}\boldsymbol{\iota}_k = k\bar{p}^k$ and $\boldsymbol{\iota}_k{}'E(\mathbf{y_k}y_i) = kp_n$, we can calculate the variances and then obtain

$$
\Delta\mathrm{E}(\pi) := \mathrm{E}(\pi_i(s_i^{R,k+1}, \mathbf{s_{-i}^{k+1}})) - \mathrm{E}(\pi_i(s_i^{N,k}, \mathbf{s_{-i}^{k}})) = \delta\left[\mathrm{Var}(s_i^{R,k+1}) - \mathrm{Var}(s_i^{N,k})\right] \sim -\varepsilon.
$$

3. $\hat{\mathbf{g}}_\mathbf{i} = \hat{\mathbf{g}} = 0$: From (2.20) to (2.22), we get

$$
\begin{aligned}
s_i^R &= a + \frac{c_H}{\tilde{d}}\tilde{t}\,y_i - c_H\frac{\varepsilon}{\tilde{d}\hat{d}}\frac{\bar{p}^n}{\bar{p}^k}\tilde{t}\,\boldsymbol{\iota}_k'\mathbf{y_k} \quad \text{and}\\
s_i^N &= a + c_H\frac{\bar{p}^{k+1}\tilde{t}}{\tilde{v}_k}y_i - c_H\frac{\varepsilon}{\tilde{d}\hat{d}}\frac{\bar{p}^n}{\bar{p}^k}\tilde{t}\left(1 - \frac{\hat{d}p_n}{\tilde{v}_k}\right)\boldsymbol{\iota}_k'\mathbf{y_k}.
\end{aligned}
$$

Noting that $\boldsymbol{\iota}'E(y_i\mathbf{y_k}) = \bar{p}^k$ and $\boldsymbol{\iota}'E(y_i\mathbf{y_k}) = kp_n$, we can calculate the variances and obtain

$$
\Delta E(\pi) \sim 2\tilde{d}\hat{d}(\bar{p}^{k+1})^2[(n-1)(p_s - p_n) + knp_n] + \bar{p}^n(\hat{d}\bar{p}^{k+1} - \tilde{v}_k)^2,
$$

where both terms are positive. ■

Chapter 3

Spatial Retail Markets With Commuting Consumers

3.1 Introduction

Since Hotelling's "Stability in Competition" (1929), economists have been concerned with the spatial distribution and the pricing policies of firms that compete oligopolistically in an output market. From the very beginning of this line of research, spatial distribution, captured by the firms' and consumers' "addresses", has also been seen as a metaphor for the differentiation of products in an abstract characteristics space.

A salient feature of most models of spatial competition is the assumption that consumers, like firms, are appropriately characterized by a single point in the address space. Moreover, the cost of travelling from a consumer's (home) location to a firm (a retail store) is added to the net price the consumer has to pay. While this may be a reasonable assumption in models where the spatial dimension refers to abstract characteristics, it can be quite unrealistic in the original context of geographical distribution. That is, for a large variety of retail markets, the relevant consumers are more appropriately described as living at one location and working at another. Commuting between these locations occurs regardless of any purchases made, and consumers will in general try to combine shopping with commuting in order to save on travel costs. In this case, the net travelling distance associated with the purchase of a good should not include any travelling along the commuting route.

That commuting matters for certain industries is most visible in the case of petrol stations, which tend to be located along radial routes of cities where commuters pass by, rather than in suburbs or in the centre. Other such examples include chain restaurants and video rental stores.

Some of the questions that arise are the following: how does commuting affect prices for given locations of firms? Here, one might conjecture that an "increase" in commuting, both in the sense of an increased proportion of commuting consumers and of an increased commuting distance, may make retail markets more competitive. Moreover, how does commuting affect location choice, both in the cases of fixed and prices and in the case of anticipated price competition? In the latter case, there are two opposite forces: more commuting might make markets more competitive and therefore induce firms not to agglomerate. On the other hand, all firms will want to be at points where the flow of commuters passing by is large, which is a force towards agglomeration. Another question is, can asymmetric equilibria emerge in this context? Can, say, the existence of high-price shops in the centre and low-price shops of the same industry in the suburbs (or the other way round) be explained by the role of commuting? Finally, to what extent do the answers to these questions depend on the precise distribution of consumers over (home, workplace)-pairs?

On a more applied level, commuting may also be an important factor in empirical studies which make use of estimates of consumers' travel costs. An example is the work of de Palma et al. (1994), who set out to predict equilibrium prices and varieties in the video cassette rental market in a Canadian city, using information on the geographic distribution of the population. In such a market, even if commuting[1] is not explicitly taken into account, one would expect that if commuting plays a role, the imputed travel costs which enter into the simulation or estimation should be lower than otherwise estimated.

In a sense, combining commuting with shopping can be interpreted as an example of a multipurpose trip, where there are economies of scale in transport. The literature on trip-chaining (cf. Eaton and Lipsey 1982, Thill and Thomas 1987, Stahl 1987) analyses the consequences of such economies associated with consumers' trips for the spatial distribution of firms, in particular, the emergence of centres. Here, in contrast, I am concerned with one market only, while the travel pattern, i.e. the commuting behaviour, is exogenously given and leads to a reduction of travel costs associated with shopping wherever the commuting and shopping routes overlap. For this reason, the relationship to the trip-chaining literature is, from a technical viewpoint, only a loose one: that literature analyses

[1] i.e. the consideration that a fair proportion of consumers will combine the renting and returning of videos with their commute

the consequences of concavities in the consumers' travel cost functions, while here the consideration of commuting makes the travel cost functions rather more convex due to the existence of flat segments for small travel distances.

It turns out that nonexistence of an equilibrium is a ubiquitous phenomenon, for both location-choice and price-competition games with commuting consumers. To illustrate how nonexistence can arise, I will in this chapter analyse in detail a particular model of price competition for given locations due to Claycombe (1991).

Claycombe (1991) introduced commuting into a model of spatial competition, arguing that for many retail markets, consumers' travel costs are in part attributable to their commuting route and not to the purchase of a good. In the model, firms and consumers are located on an infinite line. Firms are spaced at an equal distance, and consumers are uniformly distributed along the line. Each consumer commutes to a location at a constant distance from his home location. Claycombe argues that commuting makes the market inherently more competitive. In particular, he arrives at the result that if the commuting distance exceeds the distance between the firms, Nash price setting behaviour essentially leads to a competitive outcome, whereas for smaller commuting distances, commuting does not matter very much.

Claycombe and Mahan (1993) interpret this theoretical work as implying that if a population consists of both commuting and noncommuting (i.e. standard Hotelling) consumers, the equilibrium retail price depends negatively on the commuting distance of the commuters and positively on the fraction of noncommuting consumers. They go on to estimate retail prices for beef in different US cities, using commuting characteristics as explanatory variables in addition to data on concentration levels. The proportion of consumers using mass transit or car pools turns out to have a significant positive effect on prices, whereas the average commuting distance has a marginally significant negative impact. These results seem to offer support to their predictions.

My concern in this chapter is not with the econometric specification in Claycombe-Mahan linking data on prices and commuting characteristics to their predictions. Rather, I am concerned with (i) the validity of the original Claycombe (1991) analysis, and (ii) with the basis of the Claycombe-Mahan predictions, i.e. the link between Claycombe (1991) and Claycombe-Mahan (1993).

A game-theoretic analysis of a slightly simplified version of Claycombe's (1991) model shows that for small commuting distances, prices in a symmetric equilibrium depend continuously and negatively on the commuting distance and positively on the proportion of noncommuting consumers. These results both refute Claycombe's (1991) conjecture that for small distances, commuting does not matter much, and provide a more solid theoretical explanation for the predictions and empirical results of Claycombe and Mahan than their own heuristic arguments. For intermediate commuting distances, we show that a symmetric price equilibrium in pure strategies in general does not exist. Perfect competition, as predicted by Claycombe, does not result since, as the market becomes more competitive and profits decrease, firms have an incentive to cease competing for the marginal consumers and instead *increase* their price in order to extract profits from consumers over whom they have some monopoly power. Thus, while nonexistence results are fairly common in spatial models, the reason for the breakdown of equilibrium in our model is quite different from the reasons for breakdown in some other models. Only for large commuting distances, viz. at least twice the distance between firms, and only if all consumers commute, does perfect competition prevail (since in that case *every* consumer passes at least two firms on his or her commuting route). Thus, our results differ considerably from those obtained by Claycombe and at the same time are broadly consistent with the predictions and empirical results of Claycombe and Mahan (1993).

Without going into any details, I shall in the following briefly discuss two other groups of models: locational competition with fixed prices and two-stage location-then-price games.

As for locational competition, consider two or more firms locating themselves on a unit interval. Two different types of distributions of consumers are obvious candidates for the analysis: 1. distributions where there is a constant or a maximal commuting distance, and 2. a distribution such that the consumers' homes and workplaces are uncorrelated. Then roughly the following results can be obtained: 1. As long as commuting distances are not too large, an increase in the commuting distance leads to locational concentration. 2. Both large commuting distances and a large dispersion of the consumers' commuting distances are likely to give rise to nonexistence of an equilibrium for more than two firms.

To see how these results arise, notice that while in standard models the firms' location choices are partly determined by the distribution of consumers' (home) addresses, what is relevant here is rather the density of the flow of consumers passing any given point. Then some intuition for the results can be gained by drawing an analogy to results by Eaton and Lipsey (1975). An increase in commuting entails a shift of consumer density from outer points to more central points and therefore induces agglomeration of the firms. At a certain point, however, the density function may have a point of highest density, i.e. a single peak at the centre of the line. But Eaton and Lipsey (1975) have shown that there never exists an equilibrium for a number of firms exceeding twice the number of modes of the consumer density function, and we obtain a similar kind of nonexistence result here.

The case of price competition with fixed locations has been briefly discussed above and will be analysed in detail during the course of this chapter. The possible nonexistence of a price equilibrium renders the analysis of two-stage location-then-price games virtually impossible even in the duopoly case, since the second-stage cannot be evaluated for all possible subgames. As for the candidates for equilibria, the case of uncorrelated home and workplace locations is quite interesting. Here, it is clear that an equilibrium in pure strategies can never exist for symmetric locations, because then the equilibrium prices would have to be equal. But this is precisely the case where the existence of different groups of consumers described above (one indifferent, one loyal) disrupts the existence of an equilibrium. Hence, in this game, if pure-strategy equilibria exist at all, they must be strongly asymmetric: there would be a firm located near the centre that charges a high price; the other would be located further outside and charge a lower price.

3.2 Analysis of a retail market with commuting consumers

In this main section, we analyse a simplified version of Claycombe's (1991) model. First, we analyse the price equilibrium of the basic model in which all consumers commute, for a given number of firms. This model is then extended to the case in which there is a fraction of consumers who do not commute (2.2).

3.2.1 Analysis of the basic model

Firms and consumers are located on an infinite line. Along this line, firms are evenly spaced at a distance d. Firms are indexed by integer numbers, where firm i is located at $i\ d$.

Consumers are assumed to commute eastward over a constant distance c from their home location, i.e. in positive direction. In general, a consumer would be characterized by two parameters, i.e. the home and work locations. Given a constant commuting distance, however, these two parameters can be collapsed into one single parameter. For convenience, denote each consumer by the *centre* point of her commuting route, i.e. we say a consumer is located at x if she lives at $x - c/2$ and commutes to her workplace at $x + c/2$. Given this notation, let consumers be uniformly distributed along the line; more precisely, the measure of consumers is given by the Lebesgue measure.

In contrast to Claycombe, we only consider the case of constant marginal costs of production, which for simplicity are set to zero. Costs of entering the market and other setup costs are either sunk or fixed, hence are irrelevant for the price competition between the firms.

Further simplifying the Claycombe model, we assume that each consumer purchases either exactly one unit of the good or none at all.[2] The net price a consumer has to pay depends not only on the shop price of the good, but also on the travel costs associated with the purchase. If a consumer passes a shop on her commuting route, the net travel distance is zero. If the shop is located to the west of the consumer's home, the net distance is the distance between the firm and the consumer's home. The distance is analogously defined if the shop is located to the east of the consumer's place of work. Finally, travel costs are assumed to be quadratic in the net travel distance. We discuss the significance of this assumption further below. To summarise, the (indirect) utility for a consumer located at x patronising firm i to the right (east) of x is given by

$$a - p_i - t\left(\max\left\{i\; d - x - c/2, 0\right\}\right)^2 \tag{3.1}$$

Here, p_i is firm i's shop price, t is a parameter for travel cost, and a is the utility derived from the good, which we assume to be at least $(1 + c/2)^2$. [3] [4]

In the following, I only consider symmetric price equilibria in pure strategies. In the discussion following Proposition 3.1 below, however, I will address the question whether asymmetric equilibria are likely to exist.

First, let us consider commuting distances that are lower than the distance between the firms, i.e. $c < d$. In this range of values for c, each consumer passes at most one firm on her commuting route.

As a first step in deriving a symmetric Nash equilibrium in prices, we derive the demand of firm 0 charging a price p_0, assuming that all other firms charge a

[2] As will be discussed further below, this simplification does not affect our main results.

[3] This assumption merely ensures that neighbouring firms compete for customers rather than just being local monopolists, in order to make the model interesting to analyse.

[4] Compared with a model without commuting and with quadratic travel costs (e.g. d'Aspremont, Gabszewicz, and Thisse 1979), the specification (3.1) effectively introduces a flat segment in the travel cost function for travelling distances less than $c/2$. This makes more precise why formally such a model has formally not so much in common with trip-chaining models, as argued in the Introduction.

Figure 3.1: Consumer groups

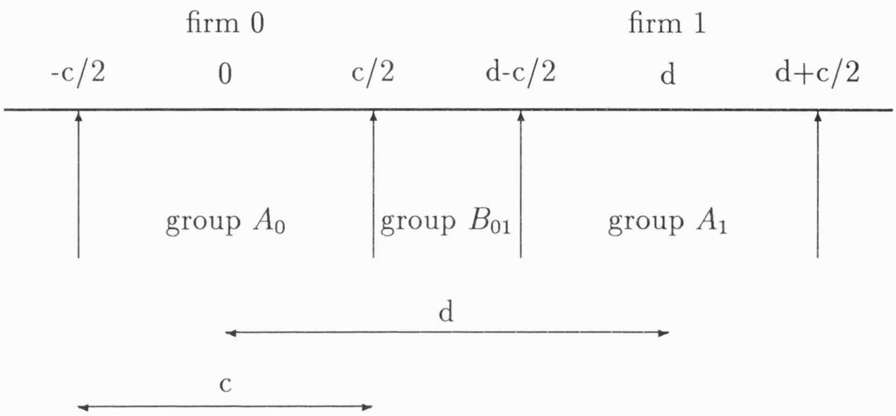

constant price p. Restricting attention to the positive (east) side of firm 0 (since the analysis is symmetric on the negative side) and to prices p_0 in a neighbourhood of p, firm 0 can potentially attract three groups of consumers (cf. Figure 3.1): Group A_0 consists of consumers who freely pass firm 0, but not firm 1. According to the notation introduced above, these are the consumers in the interval $[-c/2,\ c/2]$. Since firm (-1) charges the same price p as firm 1, only consumers in $[0,\ c/2]$ would buy at firm 1 while the consumers in the negative part would rather buy at firm (-1). A consumer located at x in $[0,\ c/2]$ decides to purchase at firm 0 as long as

$$p_0 \leq p + t\left[d - \left(x + \frac{c}{2}\right)\right]^2.$$

Defining $k := (p_0 - p)/t$, this is equivalent to $x \leq d - c/2 - \sqrt{k}$. For $k > (d - c/2)^2$, even consumer 0 will prefer to buy at firm 1 or -1, so firm 0's demand is 0. For $k \leq (d-c)^2$, on the other hand, all A_0-consumers (with total measure c) patronise firm 0. For intermediate values $k \in [(d-c)^2, (d-c/2)^2]$, finally, firm 0's demand from A_0-consumers is $2(d - c/2 - \sqrt{k})$.

Group B_{01} consists of consumers located between firms 0 and 1 who do not pass any firm on their commuting route, i.e. consumers in the interval $[c/2,\ d -$

$c/2]$. A group-B_{01} consumer located at x purchases at firm 0 if

$$p_0 + t\left(x - \frac{c}{2}\right)^2 \leq p + t\left[d - \left(x + \frac{c}{2}\right)\right]^2 \qquad \text{or} \qquad x \leq \frac{1}{2}\left(d - \frac{k}{d-c}\right).$$

Arguing similarly as above, we can show that firm 0's demand from B_{01}-consumers is $d - c - k/(d-c)$ for $k \in [-(d-c)^2, (d-c)^2]$, zero for higher prices (higher values of k), and $2(d-c)$ for lower prices.

Group A_1, finally, comprises consumers freely passing firm 1, but not firm 0, i.e. consumers in $[d - c/2,\ d + c/2]$. A consumer in this group located at x purchases at firm 0 if

$$p_0 + t\left(x - \frac{c}{2}\right)^2 \leq p \qquad \text{or} \qquad x \leq c/2 + \sqrt{-k}.$$

Demand resulting from A_1-consumers then amounts to $2[-(d-c) + \sqrt{-k}]$ for $k \in [-d, -(d-c)^2]$.[5] Adding the demand functions for the three different consumer groups gives us firm 0's total demand as a function of k:

$$D_0(k) = \begin{cases} 0 & \text{for} \quad k > \left(d - \frac{c}{2}\right)^2 \\ 2d - c - 2\sqrt{k} & \text{for} \quad k \in \left[(d-c)^2, \left(d - \frac{c}{2}\right)^2\right] \\ d - \frac{k}{d-c} & \text{for} \quad k \in \left[-(d-c)^2, (d-c)^2\right] \\ c + 2\sqrt{-k} & \text{for} \quad k \in [-d, -(d-c)^2] \end{cases} \tag{3.2}$$

This demand curve is shown in Figure 3.2 for two different values of c. In contrast to the case of no commuting ($c = 0$), which would result in a linear demand curve, commuting introduces a region of nonconcavity for higher prices. This is similar to the three-segment demand function obtained in a Linear City duopoly in which both firms each have their own hinterland with consumers who would never patronise the other firm because their reservation price is too low (cf. Gabszewicz and Thisse 1986, p.27). The nonconcavity is more pronounced the closer c approaches unity. For p to be an equilibrium price set by all firms requires

[5] (i) The lower bound for a assumed above ensures that this expression is valid, i.e. that in principle firm 0 can cover the market between $[-d - c/2, d + c/2]$ at zero price. (ii) The analysis below shows that $k < -d$ need not be considered because the requirement of nonnegative prices ensures that the relevant values always exceed $-d$.

Figure 3.2: Demand function $D_0(k)$

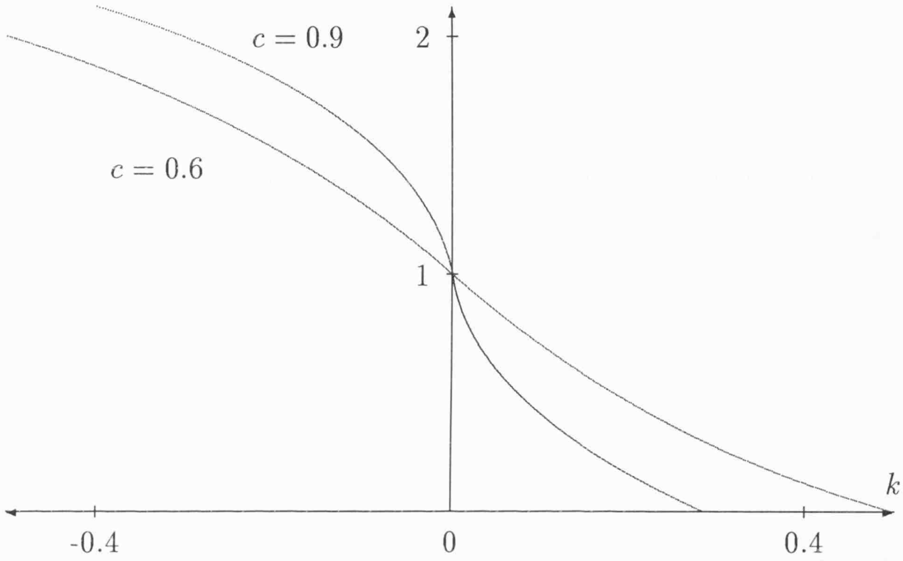

that setting $p_0 = p$ (i.e. $k = 0$) be at least locally optimal for firm 0. With firm 0's profit given by $\pi_0(k) = p_0 D_0(k)$ and $D_0(0)$ given by (3.2), the first-order condition is

$$\frac{\partial \pi_0(0)}{\partial p_0} = D_0(0) + p\frac{\partial D_0(0)}{\partial k}\frac{\partial k}{\partial p_0} = d - \frac{p}{t(d-c)} = 0, \tag{3.3}$$

hence the candidate equilibrium price is $p^* = d(d-c)t$. Our main result states the conditions under which this solution is a Nash equilibrium.

Proposition 3.1 *Let* $\bar{c} := (1/2)(13 - \sqrt{5})d \approx 0.91d$. *Then*

(a) For commuting distances $c \in [0, \bar{c}]$, *there exists a symmetric Nash equilibrium in prices in which each firm sets a (unique) price* $p^* = d(d-c)t$, *and the resulting profit is* $\pi^* = d^2(d-c)t$.

(b) For $c \in (\bar{c}, 2d)$, *an equilibrium in symmetric prices does not exist.*

(c) For $c \geq 2d$, *there is a unique symmetric Nash equilibrium in which all firms charge a zero price.*

Proof: see Appendix (Section 3.5).

Figure 3.3: Profit function $\pi_0(k)$

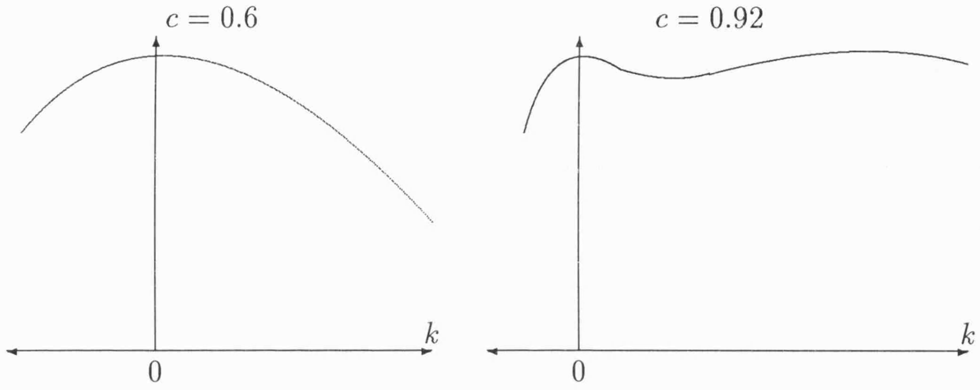

For commuting distances sufficiently smaller than the distance between the firms, prices and profits are continuous and decreasing in the commuting distance c. The reason for this is that price competition between the firms is driven by the indifference condition for the marginal consumer, who, for firms 0 and 1 and equal prices, is the consumer 1/2, a member of group B_{01}. As c increases, the net travel distance to both firms decreases. Thus, an increase in commuting distance has the same effect on price competition as a decrease in the distance of firms to each other in a standard Hotelling model.

As c approaches d, however, p^* ceases to be a Nash equilibrium (part b). This situation is illustrated in Figure 3.3, which shows firm 0's profit function for two different values of c, where all other firms charge the candidate equilibrium price $d(d-c)t$. While for values of c below $0.8d$ the profit function is concave, it becomes nonconcave for higher values. Finally, at $c \approx 0.91d$, deviating from $k = 0$ to a higher price yields a higher profit. Economically, competition for the marginal consumer drives p^* and π^* down to zero. At some point, however, a firm can profitably deviate by elevating the price to a level where only some group-A consumers close to the firm are willing to buy.

Similarly, an equilibrium does not exist if c exceeds d. The reason is the simultaneous presence of both consumers who pass two firms on their commuting route and thereby induce the firms to engage in severe price competition, and consumers who pass only one firm, implying that this firm can attract these consumers charging a positive price even if the other firms sell at zero price. This result is in striking contrast to Claycombe's (1991) prediction that for $c > d$, perfect competition prevails.

Finally, if $c > 2d$, every consumer passes at least two firms, resulting in pure Bertrand competition between neighbouring firms.

Thus, while the special case of no commuting ($c = 0$) corresponds to the case considered by Salop (1979), Proposition 3.1 shows that the existence of equilibrium can be established more generally for both $c \in [0, \bar{c})$ and $c \geq 2d$. The equilibria in these two ranges, however, are of entirely different nature.

We have shown that at most one symmetric equilibrium exists. The question remains whether other asymmetric equilibria exist, particularly in cases where a symmetric equilibrium does not exist.

It can be shown that in any asymmetric equilibrium there must be a pair of firms for which the absolute price differential exceeds $(d - c)^2$, i.e. implying that these two firms operate in the high and low segments of their demand function, respectively.

On the other hand, it can be shown that if two neighbouring firms compete only for consumers located between them, the difference in prices can in equilibrium never exceed $(d - c)^2$ because otherwise their first-order-conditions would be inconsistent. By the same argument, a situation with alternating prices, e.g. where even-numbered firms set a (the same) high price and odd-numbered firms a low price, cannot be an equilibrium if the difference between the high and low price is greater than $(d - c)^2$.

These two arguments combined suggest that for any asymmetric configuration of prices, the first-order conditions for some firms are likely to be inconsistent. Thus, while I have not been able to show *global* uniqueness of the equilibrium of Proposition 3.1 (a), these arguments lend some support to the conjecture that it is indeed unique.

In the original (1991) model, Claycombe allows for price-elastic individual demand and for increasing marginal costs. If individual demand is elastic, but marginal costs are constant, Proposition 3.1 (b) remains valid, since firms still can avoid being driven into perfect competition and hence zero profits. For rapidly increasing marginal costs, however, (short-run) perfect competition could possibly arise since at a price equal to marginal cost, firms will still be earning positive profits. In this case, it is less likely that a firm can profitably deviate.

The assumption of quadratic travel costs has been introduced into spatial models because this circumvents problems of nonexistence of a price equilibrium commonly encountered if travel costs are linear in distance (cf. Gabszewicz and Thisse 1986). Here, we have assumed quadratic costs, even though nonexistence does not arise with linear costs if there is no commuting (cf. Salop 1979). A modified model with linear costs and commuting, however, leads to results quite different from the ones obtained above, as we discuss in the following.

Assume that travel costs are linear in the net travel distance (with $t = 1$), and that firms $i \neq 0$ charge some price p. Then it can be shown that for $k \in [-(d-c), (d-c)]$, firm 0's demand is $d-k$; in particular, it does not depend on c. It follows then that the unique price in a symmetric equilibrium is $d\,t$, whatever the value of c.

To see why demand does not depend on c for intermediate values of k, assume $k > 0$ and consider the marginal consumer who is indifferent between patronising firm 0 or firm 1, and who, by assumption on k, is a member of the B_{01}-group. If her net travel distance to firms 0 and 1 is a and b, respectively, it follows that $b-a = k$ (in particular, $b > a$). An increase in c will lead to a decrease in a and b by $\Delta c/2$, leaving $b-a$ unaffected. This implies that the location of the marginal consumer does not change, and hence that demand for firm 0 remains the same.

With quadratic travel costs, on the other hand, the condition for the marginal consumer is $b^2 - a^2 = k$, and it is easy to see that the l.h.s. decreases as a and b decrease by same amounts. This consumer then prefers to purchase at firm 1. Hence the new marginal consumer must be located to the left of the previous

Figure 3.4: Profit function $\pi_0(k)$ for linear travel costs

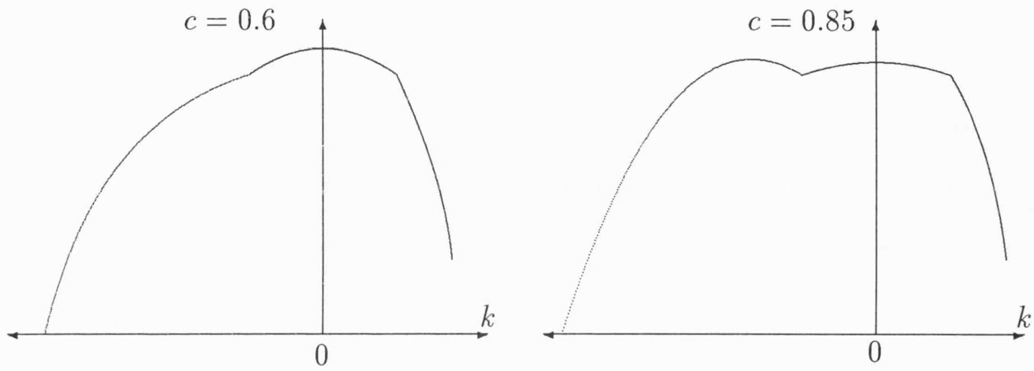

one, implying that demand for firm 0 has decreased, as reflected in the demand function (3.2).

The linear case, then, illustrates the limits of treating an increase in commuting distance as equivalent to a decrease in the distance between firms. Here, demand and equilibrium price clearly depend on the distance between the firms, but not on commuting distance. It should be pointed out, however, that this is an extreme case in the sense that any strictly convex travel cost function will lead to a demand function and an equilibrium price which qualitatively depend on the commuting distance in the same way as in the quadratic case discussed above.

A second important difference to the above analysis relates to the breakdown of equilibrium as c approaches d. It can be shown that $p^* = dt$ in the linear case ceases to be an equilibrium if c is, approximately, greater than $0.85d$, because firm 0 will have an incentive to *undercut* its rivals. This situation then, shown in Figure 3.4, is rather similar to the reason for nonexistence in a standard Hotelling model with linear travel costs if firms are close together.

3.2.2 Commuting and noncommuting consumers

We modify the model of the previous subsection by assuming that a proportion α of the consumers does not commute. The significance of this assumption is the fact that a proportion of consumers uses mass transit (or car pools) to commute rather than their own car. These consumers, then, are less likely to combine purchases with their commuting, as pointed out by Claycombe and Mahan (1993). Thus, if their travel costs for shopping are associated with the purchases to full extent, this is appropriately modelled by the assumption that these consumers do not commute at all.

Adhering to the notation of the Section 2.1, a noncommuting consumer x located between firms 0 and 1 will patronise firm 0 if

$$p_0 + tx^2 \leq p + (d - x)^2 \qquad \text{or} \qquad k \leq d - 2x.$$

Hence, in the range $k \in [-d, d]$, firm 0's demand from noncommuting consumers is $\alpha(d - k)$. Combining this with the demand by commuting consumers given by (3.2), we obtain the total demand for firm 0 for small absolute values of k:

$$\begin{aligned} D_0(k) &= (1-\alpha)\left(d - \frac{k}{d-c}\right) + \alpha\,(d-k) = d - k\left(\frac{\alpha}{d} + \frac{1-\alpha}{d-c}\right) \\ &\text{for} \quad k \in [-(d-c)^2, (d-c)^2]. \end{aligned}$$

Proceeding similarly as in Section 2.1, we can derive the price and profit for the unique candidate for a symmetric equilibrium.

Proposition 3.2 *For any given $\alpha < 1$ there exists a $\bar{c}(\alpha) < d$ such that*

(a) For $c \in [0, \bar{c}(\alpha)]$, a symmetric price equilibrium is given by

$$p^* = d^2 \frac{(d-c)t}{d - \alpha c} \qquad \textit{and} \qquad \pi^* = \frac{d^3(d-c)t}{d - \alpha c}. \tag{3.4}$$

(b) For any $c > \bar{c}(\alpha)$, a symmetric equilibrium does not exist.

Proof: see Appendix (Section 3.5).

Similarly as before, $\partial p^*/\partial c < 0$, thus prices fall as the commuting distance increases, where the magnitude of the effect varies proportionally with the fraction of commuting consumers. Moreover, $\partial p^*/\partial \alpha > 0$, thus prices decrease continuously as the proportion of commuting consumers becomes larger. This result nicely corresponds to Claycombe and Mahan's (1993) finding that the proportion of consumers using mass transit or car pools has a significant positive effect on retail beef prices.

For commuting distances close to or larger than d, however, the nonexistence problem encountered in the basic model still arises. According to (3.4), even in the presence of noncommuting consumers, firms still compete for the marginal commuting consumer, which drives prices and profits to zero as c approaches d. Then, again, for some critical value of c, a firm can profitably deviate from p^* by extracting a positive price from its local consumers.

Finally, even for large commuting distances ($c > 2d$), perfect competition never prevails since positive profits can be earned from the fraction α of noncommuting consumers.

How does the critical value $\bar{c}$ depend on α? Unfortunately, a closed expression for $\bar{c}(\alpha)$, which is a non-monotonic function, cannot be obtained for this generalised model. Based an numerical experiments, we make the following observation:

Observation 3.1 *(a) There is a lower bound to $\bar{c}$ at around $0.85d$, i.e. for any value of α, p^* as given in (3.4) is always an equilibrium price as long as $c < 0.85d$. The lowest values of $\bar{c}$ are obtained for values of α of around 0.7.*

(b) As α approaches unity, $\bar{c} \rightarrow d$.

Therefore, whatever the value of α, c has to be fairly large for nonexistence to obtain. Moreover, since for $\alpha \rightarrow 1$ the model converges to a standard model without commuting, there is no upper bound to c less than d. Nevertheless, any positive fraction of consumers who do commute leads to the breakdown of equilibrium for every $c > \bar{c}$ (Proposition 3.2).

3.3 A review of Claycombe's analysis and the predictions of Claycombe-Mahan

In this section we address the concern mentioned in the Introduction, viz. the validity of Claycombe's (1991) analysis and the link between this work and Claycombea and Mahan (1993).

As to the latter, we observe that the Claycombe-Mahan predictions cannot be derived from Claycombe (1991). According to that analysis, perfect competition prevails whenever $c > d$, whether there is a fraction of noncommuting consumers or not. In a free-entry equilibrium, therefore, this situation can never arise, since firms would be making losses. However, the paper remains silent on the properties of the price equilibrium for the case $c < d$.

On the other hand, the theoretical explanations Claycombe and Mahan provide for their predictions are only heuristic and not directly related to the Claycombe (1991) model: average commuting distance determines the number of shops consumers pass on average, which in turn partially determines the breadth and hence competitiveness of (local) markets, and therefore, prices. Moreover, consumers using mass transit or car pools "are unlikely to use the commute to stop at a store...Hence, a high proportion of these consumers is expected to generate high prices".

At this point it is appropriate to return to Claycombe's (1991) analysis in order to see where the difficulties arise.

On p.308, Claycombe correctly argues that for $c > d$ and a given number of firms, price shading will occur as long as price exceeds marginal cost. On the other hand, the "prediction of competitive pricing" which seems to follow from this is valid only if competitive pricing can actually be sustained as a Nash equilibrium. Our Proposition 3.1 shows that this is not the case.

Claycombe then discusses the free-entry equilibrium for a situation with some fixed cost and a constant marginal cost. With this cost structure, competitive pricing implies losses for the firms, therefore they exit (or more appropriately, they do not enter in the first place). Assuming that the distance between the firms is inversely related to the number of firms, as is the case e.g. if firms are spaced evenly on a circle rather than a line, this line of reasoning implies that under the free-entry assumption there can never be an equilibrium for which $c > d$. However, this does not tell us anything about the actual equilibrium outcome. There is no reason to assume that "when exit drives d up to the level of c an equilibrium is reached" (p.309) in which prices can be held above marginal cost, even if an equilibrium existed (contrary to Proposition 3.1 b). On the contrary, by continuity, Proposition 3.1 (a) would imply that for $c = d$, competitive pricing would still prevail, leading to losses for the firms. If, on the other hand, supramarginal prices are merely due to an integer number of firms and the required nonnegativity of profits, this result is quite obvious and has nothing to do with commuting.

Claycombe's above result that for $c > d$, perfect competition prevails, is carried over to the case where there is a fraction of noncommuting consumers (mixed demand, p. 310). According to our analysis (Section 2.2), however, in this case a symmetric pure-strategy equilibrium does not exist even if $c > 2d$.

For small commuting distances ($c < d$), finally, Claycombe is concerned with the derivation of market boundaries for given prices, but does not carry this analysis further to determine equilibrium prices. The results of the analysis here, as given by Proposition 3.1, certainly are at odds with Claycombe's assertion that for this parameter range "the model not differ substantially from the model where there is no commuting at all" (p. 310). At the same, they provide a more rigorous theoretical basis for the predictions of Claycombe and Mahan, which appear to be supported by their data.

3.4 Concluding remarks

In this chapter I have shown that introducing commuting into an otherwise standard spatial model has a significant impact on the price competition among retail firms. For commuting distances which are small compared to the distance between firms, prices are decreasing in both the commuting distance and the proportion of commuting consumers. For larger commuting distances, however, a price equilibrium in general does not exist. Only in the extreme case where *each* consumer freely passes at least two shops on his or her commuting route (i.e, in particular, all consumers commute), is perfect competition the equilibrium outcome. Our results thus strikingly differ from those obtained by Claycombe (1991).

The results of our model appear to be, at first glance, similar to results frequently encountered in spatial models: (i) As firms move closer to each other (directly or indirectly), equilibrium prices decrease. (ii) On the other hand, nonexistence is a common result particularly if firms are close to each other (cf. Gabszewicz and Thisse 1986).

A closer look, however, reveals some deviations from this pattern. First, the analogy between an increase in commuting distance and a decrease in the distance between firms has its limits, as the discussion of the case of linear travel costs has shown. Second, and more importantly, the reason why a price equilibrium ceases to exist for longer commuting distances is very different in nature from the breakdown of equilibrium in the standard Hotelling model with linear travel costs. There, if firms are close together, a firm will be tempted to shade its price in order to capture a large share of the market of the neighbouring firm; here, in contrast, a firm will give up competing for the marginal consumer located between the firms and instead *increase* the price in order to extract profits from its local consumers.

The nonexistence problem is a quite general phenomenon which can be shown to arise in other spatial models with commuting consumers as well, e.g. models in which the commuting distance is not the same for every consumer. It arises because of the presence of both consumers who freely pass two shops on their commuting route and thus are only concerned about the shop price, and other consumers who pass only one shop and hence have to incur travel costs if they purchase at another shop. Prices above marginal costs cannot be sustained as firms compete for consumers in the former group, whereas on the other hand firms can always secure positive profits by charging supramarginal prices from consumers in the latter group.

Thus, while the assumption of a constant commuting distance in this chapter is quite restrictive, relaxing it only makes nonexistence problems more likely to occur. Consumers commuting over long distances make the market more competitive, whereas the existence of local consumers who commute only over short distances (or not at all, cf. 3.2.2) gives firms some monopoly power, disrupting competiton for the marginal consumer.

Theoretical and empirical work both suggest that commuting plays an important role for the price determination in retail markets, and commuting is also likely to affect firms' location choices as well. The nonexistence of pure-strategy price equilibria in apparently plausible models, however, poses a serious obstacle to the analysis of models involving location choice. Further research will hopefully lead to models in which these issues can be analysed.

3.5 Appendix: Proofs

Proof of Proposition 3.1:

1. First consider the case $c < d$. Obviously, since (3.3) is a necessary condition for a symmetric price equilibrium, $p^* = d(d-c)t$ is the only candidate for an equilibrium price. Given the concavity of π_0 for $k \in [-(d-c)^2, (d-c)^2]$, p^* is locally optimal in this price range. It remains to be shown under which conditions deviating to a price outside this range will not be profitable.

1.1 Consider a deviation by firm 0 from p^* such that $k > (d-c)^2$. Since $p_0 = kt + p^*$, by (3.2) firm 0's profit is

$$2t[k + d(d-c)t](d - c/2 - \sqrt{k}).$$

Therefore, deviating is profitable if

$$2[k + d(d-c)t](d - c/2 - \sqrt{k}) - d^2(d-c) > 0. \tag{3.5}$$

In order to obtain the maximum of this expression in k, we take a look at the first-order condition

$$d - c/2 - \sqrt{k} = \frac{1}{2\sqrt{k}}[k + d(d-c)t]. \tag{3.6}$$

Unless c is relatively close to d, this equation has no real root in k, in which case the l.h.s. of (3.5) is decreasing in the entire relevant range. For larger c, the solution to (3.6) is

$$\sqrt{k} = \frac{1}{6}\left(2d - c + \sqrt{c^2 + 8cd - 8d^2}\right), \tag{3.7}$$

where indeed a maximum is obtained. Substituting this expression in (3.5) and furthermore expressing c as a fraction of d, $c = \lambda d$, the l.h.s. of (3.5) becomes

$$2\left(1 - \frac{\lambda}{2} - \frac{2-\lambda}{6} - \frac{1}{6}\sqrt{\lambda^2 + 8\lambda - 8}\right)\left[1 - \lambda + \left(\frac{2-\lambda}{6} + \frac{1}{6}\sqrt{\lambda^2 + 8\lambda - 8}\right)^2\right]$$
$$-(1-\lambda). \tag{3.8}$$

In the relevant range for λ, i.e. between $\sqrt{24} - 4$ and one[6], (3.8) is monotonically increasing and has a unique root at $(13 - 5\sqrt{5})/2$, which gives us the critical value

[6] For smaller λ, (3.7) is not real.

$\bar{c}$ of the Proposition[7]. We have thus shown that a deviation from p^* to a higher price such that $k > (d-c)^2$ is not profitable if and only if $c \leq \bar{c}$.

1.2 To complete part (a), we have to show that deviations below p^* such that $k < -(d-c)^2$ are not profitable either. First of all, since p_0 cannot fall below 0, the relevant range for k is $[-(d-c), -(d-c)^2]$. Similarly as above, using (3.2) the condition for a profitable deviation is

$$[k + d(d-c)](c + 2\sqrt{-k}) - d^2(d-c) > 0.$$

This expression has its maximum at $k = -(d-c)^2$, where the value is $-(d-c)^2$. Therefore, such a deviation is never profitable.

2. Now let $c \in (d, 2d)$.

(a) Then there exists a group of consumers located between $[d - c/2, c/2]$ who can reach both firms 0 and 1 at zero cost. A price $p > 0$ charged uniformly by all firms is not sustainable as an equilibrium since a firm could, by shading its price by an infinitesimal amount, increase its demand by $2(c-d)$ and thus increase profits.

(b) On the other hand, $p = 0$ is not sustainable either because consumers located in $(-d + c/2, d - c/2)$ reach firm 0 at zero cost, but neither of the neighbouring firms. By charging an arbitrarily small price, firm 0 can thus attract a subset of this group and thereby attain positive profits. Therefore, a symmetric price equilibrium does not exist. By continuity, arguments 1. and 2. extend to the case $c = d$.

3. Finally, for $c \geq 2d$, argument 2a still applies, but not argument 2b, as *every* consumer can reach at least two firms. Therefore, $p^* = 0$ is the unique symmetric equilibrium. ■

Proof of Proposition 3.2:

1. First of all, for p to be a symmetric equilibrium price (implying $k = 0$)

[7] This result was obtained with the help of *Mathematica*.

requires

$$\frac{\partial \pi_0(0)}{\partial p_0} = d - \frac{p}{t}\left(\frac{\alpha}{d} + \frac{1-\alpha}{d-c}\right) = 0,$$

which leads to the candidate equilibrium (3.4).

2. To show the existence of a critical value $\bar{c}(\alpha)$, we first have to show that for any $\alpha < 1$, there exists c such that a deviation from p^* into to the range $k \in [(d-c)^2, (d-c/2)^2]$ is profitable, which (proceeding as in the proof of Proposition 3.1, cf. (3.5)) will be the case if

$$\begin{aligned} r(\alpha, c, k) \;&:=\; \left[(1-\alpha)(d-c-2\sqrt{k}) + \alpha(d-k)\right]\left(k + \frac{d^2(d-c)t}{d-\alpha c}\right) \\ &\quad -\frac{d^3(d-c)t}{d-\alpha c} > 0. \end{aligned} \tag{3.9}$$

Specifically, consider $\bar{k} := (d-0.75c)^2$. Then $r(\alpha, c, \bar{k}(c) > 0$ is equivalent to

$$s(\alpha, c) := 256(1-\alpha c) r(\alpha, c, \bar{k}(c)) > 0,$$

where s, a polynomial of fifth order in c, is continuous in both parameters and strictly increasing in c for all values of α. Since $s(\alpha, 1)$ is unambiguously positive, continuity and monotonicity of s in c imply the existence of a $\tilde{c}$ such that $s(\alpha, c) > 0$ for all $c \in [\tilde{c}, 1]$ and $s(\alpha, c) > 0$ for all $c < \tilde{c}$. By choosing k to maximise the l.h.s. of (3.9) one then obtains the smallest value $\bar{c}$ such that this holds. ■

Chapter 4

Product Differentiation, Uncertainty and the Stability of Collusion

4.1 Introduction

Does heterogeneity of products limit the scope for collusion among oligopolists, or is it rather a facilitating factor? Many economists would hold the former to be true, arguing that heterogeneity, in some sense, makes the firms' coordination problem more complex. A recent surge of interest in this question[1], however, has led to numerous game-theoretic models which seem to suggest the interpretation that cartel stability *increases* with the degree of product differentiation.

In contrast, this chapter argues that uncertainty, neglected in both the informal literature and the new theoretical contributions, alters the problem of sustaining collusion in a fundamental way and plays in a crucial role in determining the effect of product differentiation on the scope for collusion. I introduce the idea that an increase in the heterogeneity of products leads to a decrease in the correlation of the demand functions for the goods. In an environment where a firm cannot observe its rivals' actions but has to infer from observable signals whether another firm has deviated, this in turn makes collusion more difficult to sustain, as discriminating between random demand shocks and deviations from the cartel strategy becomes more difficult.

A first simple model shows how uncertainty can give rise to demand functions the correlation of which depends on the degree of product differentiation. Two other duopoly models quite different from each other illustrate how the correlation of demand functions is linked to the stability of collusion, and generate similar results: given a sufficiently high discount factor and a sufficient level of demand uncertainty, collusion becomes *less* sustainable as products become more differentiated, and if product differentiation exceeds a critical level, collusion may not

[1] For horizontal product differentiation focused on here, this includes the papers by Deneckere (1983), Wernerfelt (1989), Chang (1991), Ross (1992) and Häckner (1993), and the discussion in Martin (1993, p. 116). The case of vertical product differentiation is studied by Häckner (1994). In contrast, the papers by Jehiel (1992) and Friedman and Thisse (1993) study the effect of collusive conduct on firms' effort to differentiate products, and Zhang (1995) can also be counted to this strand of literature.

be profitable at all. These results stand in sharp contrast to those in the recent theoretical literature. At the same time, the model provides a simple analytical foundation for the traditional view that heterogeneity limits the scope for collusion.

According to the traditional view, heterogeneity impedes cartel behaviour because, in some sense, firms face a situation of "higher complexity". For example, while with homogeneous products firms merely have to agree on one price, with heterogeneous goods a whole array of prices has to be negotiated. This problem "grows in complexity by leaps and bounds" (Scherer and Ross 1990, p. 279) with the number of chararacteristics in which the goods can differ. Similarly, firms may have difficulty in monitoring the policies of their rivals in complex situations (Clarke 1985, p. 60).[2]

Though intuitively compelling, it is difficult to pin down analytically an appropriate interpretation of this argument. For example, one could argue that with heterogeneous goods, the relevant space of the products' attributes becomes very large and may not even be specifiable in advance, and thus would render both cartel negotiations and subsequent enforcement increasingly difficult. However, in standard models of product differentiation, products are usually symmetrically positioned in a relatively simple space of characteristics. Here, it is harder to see in which sense differentiation could lead to a situation of increased complexity.[3] Put differently, if the intuition that heterogeneity has something to do with complexity is correct, then product differentiation has implications not captured by the standard models.

[2] Posner (1976, p.60) is more precise, but also makes specific explicit assumptions about the information structure and the nature of cartel agreements: "... the detection of cheating by members of the cartel will be complicated by the difficulty of knowing whether a competitor's price is below the agreed level or is simply a lower price for a lower grade or quality of the product."

[3] Similarly, Tirole (1988, p.240) notes that while the role of detection lags as a factor hindering collusion is well understood, "efforts to formulate the second factor [asymmetries, as which he counts the case of differentiated products] have not been as successful".

Empirical evidence on this issue suffers from measurement problems of two kinds, the measurement of product heterogeneity, and the verification of collusive behavior. As far as products are concerned, the most frequent method of assessing the degree of differentiation is simply by inspection. Cases of collusive behaviour, on the other hand, are often drawn from antitrust sources.[4] Here, the existing support for the traditional view homogeneity eases collusion of course raises the question of whether homogeneity of products may sometimes be a reason or condition for firms being indicted in the first place.[5]

A different and perhaps more convincing kind of support for the traditional view can be found in case studies which emphasize the role of strategic product *standardization*. For example, in his study of the US electrical industry, Sultan (1974, p.28-29) points out that the main purpose of an "organized industry effort to standardize the designs of most products" during the 1920s was to reduce price warfare by increasing the visibility of price-cutting.

From these observations we may conclude that the available evidence for the traditional view is rather weak. On the other hand, I am not aware of any evidence supporting the opposite conjecture that product differentiation is a facilitating factor.[6]

The traditional view has been challenged by theorists who pointed out that with differentiated goods, both the benefit of collusion, and the gain of deviating from a cartel strategy, are likely to be smaller than in the case of homogeneous goods. Hence, a priori little can be said about the net effect of product differentiation on the stability of collusion.

Following Deneckere (1983), several researchers have recently analysed the ef-

[4]cf. Hay and Kelley (1974), Fraas and Greer (1977).

[5]In fact, Hay and Kelley (1974) find that "virtually all of the entries [in their sample] would read 'high' product homogeneity."

[6] This assertion is consistent with an important observation made by Fraas and Greer (1977): where the environment is most conducive to collusion, firms can collude tacitly, whereas under less favourable conditions firms may require formal cartel arrangements to sustain collusion. In this sense, collusion will tend to be most visible when colluding is more difficult.

fect of product differentiation on cartel stability within game-theoretic models. All models that have been studied are deterministic, models in which deviations from a cartel strategy are detected immediately and precipitate retaliation. Therefore, sustainability of collusion only depends on the tradeoff between the benefit from collusion, and the gain obtained by deviating from a cartel strategy.

Cartel stability is measured by the critical discount rate below which the joint profit maximising price can be sustained with a trigger strategy (in most cases a simple grim trigger strategy; for an exception, see Wernerfelt [1989]). For this case, the critical discount rate is simply the ratio of the collusive benefit to the defection gain (cf. Martin 1993, p.104). A decrease in the critical discount rate is then interpreted as a decrease in the scope for collusion.

From the argument sketched above, there would be little reason to expect any systematic relationship between the critical discount rate and the degree of product differentiation, and indeed there are several results showing an ambiguous relationship.[7] Nevertheless, there exists a variety of models in which the critical discount rate *increases* as products become more differentiated. This has led some to question the validity of the traditional view (cf. Ross 1992), and has even led to the emergence of a new conventional wisdom among theorists.

A serious problem with this research programme is the disregard of any uncertainty which might play a role. In particular, it is implicitly assumed that any deviation from collusive behaviour is detected with certainty. With this information structure, firms can only choose between two extremes: either to adhere to the cartel strategy, or to cut the price (or increase the quantity) by a large amount so as to maximise the current-period profit, in anticipation that retaliation will follow with certainty. In contrast, one would more realistically expect firms to consider increasing their profits by cutting their price only *slightly*, in the hope that such a deviation will go unnoticed by the other firms. A formal analysis of such considerations, however, which is one of the main purposes of this chapter, requires a framework with imperfect monitoring.

[7] E.g. Deneckere 1983, Ross 1992, Wernerfelt 1989.

In this chapter, I add a new dimension to the analysis of cartel stability, viz. uncertainty and, hence, the *probability* of price wars being triggered in the first place. Specifically, while in the case of completely homogeneous goods firms face the same demand function, it seems reasonable to assume that if demand is stochastic, the correlation of the demands for the products is likely to be lower the more heterogeneous these goods are. Intuitively, the more similar two goods are, the more similar are the customers attracted by these goods. Shocks on the demand side, affecting entire groups of customers, will then be reflected in these firms' demand functions in largely the same way. On the other hand, the more differentiated two goods are, the lower the correlation of demand is likely to be because the the firms attract customers from largely distinct populations. In an environment in which firms cannot observe their competitors' actions but have to infer from publicly observable noisy signals whether anybody has deviated from the cartel strategy, this correlation effect then decreases the scope for collusion, since discriminating between cheating and exogenous random fluctuations in demand is more difficult.

Three specific models serve to illustrate these ideas. In Section 4.2, it is shown in a Hotelling-type duopoly model how a combination of random "macro" shocks on the demand side and heterogeneity in tastes among consumers generate positively correlated demand functions for the two goods, such that the correlation coefficient depends positively on the degree of substitutability between the goods. The model, therefore, provides an analytical foundation of the idea suggested above. In this first model, the correlation is derived endogenously. In two models analysed subsequently which are more convenient vehicles in which to analyse collusion, demand correlation is simply assumed to vary with product substitutability.

Section 4.3 sets out the general framework for collusion. Two specific examples are discussed in Sections 4.4 and 4.5. These are duopoly models with price-setting firms and demand uncertainty. The demand shocks belong to a specific family of distributions with compact support, and their correlation depends positively on the degree of substitutability. Each firm can observe its own and the rival's realised demand, but cannot observe either the shocks or the other firm's price. Firms collude using trigger strategies of a generalized Green-Porter (1984) type. During a collusive phase, firms set a certain cartel price. The observation of certain realised demand vectors triggers a "price war". During such wars, firms play the static Bertrand equilibrium for a fixed number of periods, and thereafter return to the collusive mode. While in general the structure of such strategies can be very complicated, we show that for the family of distributions considered here, collusive strategies which are payoff-maximising within this class have a particularly simple structure.

The sustainability of a certain cartel price requires that two incentive constraints be satisfied, neither of which is implied by the other. First, it must not be profitable to undercut the price by a large amount in order to maximise the current-period payoff, if it is anticipated that retaliation will follow with certainty. This is the familiar constraint analysed in the theoretical literature discussed above. Second, it must not be profitable to cut the price marginally, i.e. the marginal gain of doing so must be counterbalanced by a loss due to an increase of the probability of a price war being triggered. This sort of constraint is familiar from Green and Porter (1984) and related works. We point out that due to a fundamental and quite general nonconcavity in the payoff function, one constraint does not imply the other, hence both constraints have to be taken into account.

Two specific examples illustrate how the sustainability of collusion depends on the degree of substitutability. The first (Section 4.4) is a Hotelling-type model. The second (Section 4.5) is a model with a linear demand system derived from a representative consumer's utility function. While these models are very different in terms of how the Bertrand and monopoly prices depend on the degree of substitutability, the analysis of collusion leads, apart from differences in the details, to quite similar results:

1. While sustainability of a certain cartel price is by and large negatively

related to the substitutability of the goods as far as the incentive constraint for *large* price cuts is concerned (as already shown in other papers), this relationship is reversed if the constraint for *marginal* cuts is the binding constraint.

2. For lower (actual) discount factors, both constraints are relevant, and this can lead to a nonmonotonic relationship. A certain price may not be sustainable for homogeneous goods or for a large degree of differentiation, but only for intermediate degrees of differentiation. For higher discount factors, however, the constraint for marginal price cuts remains as the only relevant constraint (provided there is sufficient demand variation, in the case of the Hotelling model). Here, the predictions are in line with the traditional view and stand in contrast to those of the recent theoretical literature: a certain price may be sustainable for sufficiently homogeneous products, but not for larger degrees of differentiation. Moreover, if differentiation exceeds a critical level, collusion (at whatever price) may not be profitable at all.

Thus the similarity of the general picture suggests that a reasonably robust economic mechanism is at work here. While the equilibrium probability of a price war (which happens to be zero in the models studied here) affects a firm's discounted payoff, it is the *increase* of this probability in case of a deviation that matters for the sustainability of a certain strategy. This is why the correlation of demand is important: the more differentiated the products are, the lower is the correlation, and therefore the smaller is the effect of a deviation on the probability of a price war.[8] So the retaliation phase loses its deterrent effect, which undermines the stability of the cartel.[9]

[8] Depending on the particular oligopoly model, this effect might be reinforced by a lower cross-price elasticity of demand functions.

[9] More precisely, while the set of events that trigger a price war is itself endogeneous, adjusting this set cannot cancel the described effect. Rather, cartel stability is undermined because with more heterogeneity, optimally designed strategies must be more lenient than in the case of homogeneous goods.

4.2 Demand shocks and correlated demand functions

This section formalises the idea discussed in the Introduction, that in a market with demand uncertainty, the correlation of demand shocks is likely to be higher for rather similar goods than for more differentiated goods. The model used here is simply a modified version of the well-known Hotelling model with taste heterogeneity due to de Palma et al. (1985).

Two firms are located symmetrically on a line $[0, 1]$. With $\sigma \in [0, 1]$, firm 1 is located at $\sigma/2$ and firm 2 at $1 - \sigma/2$. Thus σ measures the degree of product substitutability: for $\sigma = 1$, we have homogeneous goods, and for $\sigma = 0$, maximal differentiation.

Consumers are distributed along the line. The utility of a consumer located at z that purchases firm i's product is given by

$$w_i(z) = y + a - p_i - b(z - z_i)^2 + \varepsilon_i,$$

where y denotes income, a is the utility derived from consuming the most preferred good, b is a parameter for "travel costs", p_i is firm i's price and z_i its location. Heterogeneity of tastes is introduced by means of the random variables ε_1 and ε_2, which are assumed to be i.i.d. double exponentially distributed with zero mean and variance $\mu^2\pi^2/6$ (cf. Anderson, de Palma and Thisse 1992, p.363). In effect, through this random utility specification a second dimension of product differentiation is introduced into the model, in addition to differentiation along the Hotelling line.[10] The degree of differentiation in this second dimension is measured by the parameter μ. With the distribution of ε_i as specified above, the probability that consumer z buys at firm 1 is then given by

$$P_1(z) = \left\{1 + \exp\left[\frac{1}{\mu}\left(p_1 - p_2 + (1 - \sigma)(2z - 1)\right)\right]\right\}^{-1}. \qquad (4.1)$$

[10] For an extensive discussion of this formalisation of product differentiation, see Anderson, de Palma and Thisse (1992).

In contrast to de Palma et al. (1985) and other models, I assume there are "macro" shocks affecting the density of consumers in different parts of the Hotelling line. More specifically, the consumers fall into two groups, group 1 in the interval $[0,1/2)$ and group 2 in $(1/2,1]$. The densities of consumers in both intervals, u_1 and u_2, are independent and uniformly distributed over $[0,2]$.[11] These density shocks can be thought of as being caused by taste changes, business cycles, or other reasons, affecting groups 1 and 2 in different ways. Other things equal, consumers in group 1 have a preference for good 1, and similarly for group 2. But with taste heterogeneity, there is always a positive probability that a consumer will purchase the "other" product, where this probability depends both on prices and the similarity of the products. This is obviously a very crude and simple way of introducing uncertainty into the model; many other specifications are conceivable.

The description of demand has two salient features. First, the random utility approach used here implies a "cross-over" of market areas[12]; i.e., even with differentiation, both firms attract consumers from the entire Hotelling line. Second, the density shocks (i.e. market size shocks) in both halves of the line imply that the correlation of the firms' demands depends on the proportions of customers each firm draws from each half of the line, and therefore depends on the degree of product differentiation.

Firm 1's demand is obtained by integrating the purchase probabilities (4.1) over the entire line, taking the densities u_1 and u_2 into account:

$$q_1 = u_1 \int_0^{1/2} P_1(z)\mathrm{d}z + u_2 \int_{1/2}^{1} P_1(z)\mathrm{d}z \tag{4.2}$$

and firm 2's demand is obtained analogously. Here, we are only concerned with the correlation of the demand functions for the case where both firms charge equal prices. Using (4.2) and the corresponding expression for firm 2, with $p_1 = p_2$ we

[11] The assumption of independence simplifies calculations but is not necessary; any degree of positive correlation would leave the qualitative features of the results unaffected.

[12] Cf. Archibald et al. (1986) and de Palma et al. (1994).

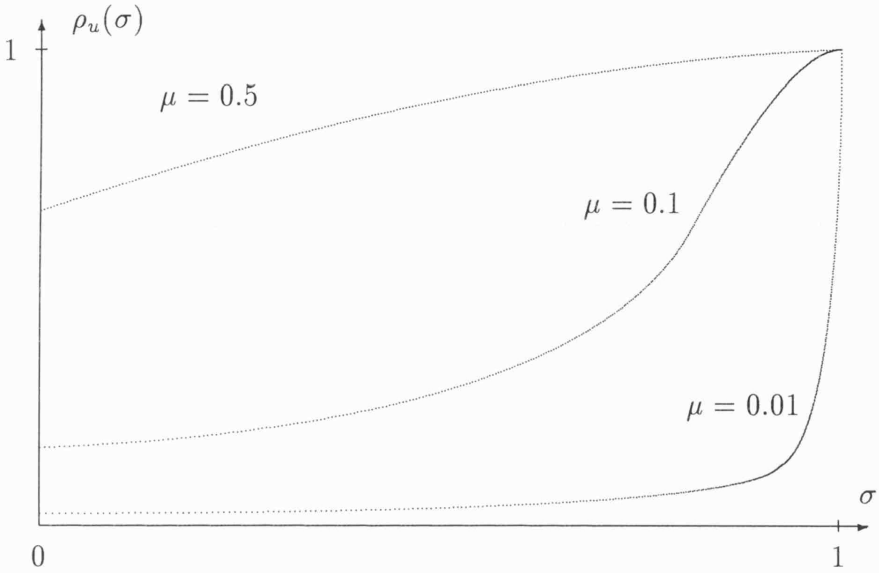

Figure 4.1: Correlation ρ_u between q_1 and q_2.

have $q_1 = u_1 m_1 + u_2 m_2$ and $q_2 = u_1 m_2 + u_2 m_1$, where

$$m_1 = \frac{1}{2} - \frac{\mu}{2(1-\sigma)} \log \frac{2}{1 + e^{-\frac{1-\sigma}{\mu}}} \quad \text{and} \quad m_2 = \frac{1}{2} - \frac{\mu}{2(1-\sigma)} \log \frac{1 + e^{\frac{1-\sigma}{\mu}}}{2}.$$

Since u_1 and u_2 are independent, the correlation between q_1 and q_2 is then given by $\rho_u = 2m_1 m_2/(m_1^2 + m_2^2)$. How this correlation depends on the substitutability parameter σ is depicted in Figure 4.1. For σ close to one, the correlation is close to one as well: since the products are vitually identical, consumers buy either of the two goods with almost equal probability. On the aggregate level, the demand functions become more or less identical (i.e., both m_1 and m_2 are close to $1/2$). In contrast, for larger degrees of differentiation, group-1 consumers will buy good 1 with a higher probability than good 2, and vice versa for group-2 consumers. On the aggregate level, this implies that firm 1 will draw its customers to a larger extent from group 1 than from group 2 (formally, $m_1 > m_2$), and therefore will be more affected by the shock m_1 than by m_2. Since the reverse holds for firm 2, this implies a lower correlation between the two firms' demand functions. This is precisely the idea dicussed less formally in the Introduction.

The correlation also depends on the heterogeneity parameter μ: for small μ, i.e. little heterogeneity, even slight differentiation along the Hotelling line will suffice to effectively separate the firms' market areas into the two halves of the line. In contrast, with large μ, the correlation of the firms' demand functions will be considerable even with maximal differentiation along the Hotelling line.

A second effect in this model is that with heterogeneous goods, the variance of each firm's demand is larger than with more similar goods. Rather than being an artifact of the particular distributions considered here, this relationship seems economically very plausible: a firm that targets a specific group of customers is mainly affected by variations of demand by that particular group, whereas a firm that offers a standard product purchased by different customers is less vulnerable to shocks affecting a particular group.

While this model is very convenient for showing how correlated demand functions can be derived, the resulting expressions are awkward to handle when it comes to the analysis of collusion. In what follows, therefore, I will study models with somewhat simpler demand functions. Moreover, the dependence of the correlation of q_1 and q_2 on σ will be assumed rather than derived as here. Thus, this section should be seen as providing the motivation and justification for what constitutes a key assumption in this chapter.

4.3 Collusion: general framework

This section sets out the general framework of the specific models that will be discussed in Sections 4.4 and 4.5. In order to obtain analytically tractable models, a number of restrictive assumptions have to be made. Thus, "general" here refers to those parts of the models and the analysis that are common to the two examples discussed below, rather than to any truly general model.

4.3.1 Basic model and information structure

We study collusion in a duopoly with price-setting firms and differentiated products. The substitutability of the goods is characterized by the parameter $\sigma \in [0,1]$, where for $\sigma = 1$ the goods are homogeneous, and for $\sigma = 0$, maximally differentiated (though not necessarily independent).

Demand uncertainty is introduced through two random variables u_1 and u_2, and in the most general formulation, both firms' realised demands, q_1 and q_2, are functions of both the price vector $\mathbf{p} = (p_1, p_2)$ and the demand shock $\mathbf{u} = (u_1, u_2)$.

I assume that while a firm cannot observe its rival's price, it can observe both its own and the rival's realised demand.

The assumption of cartel models in the tradition of Stigler (1964) that firms cannot observe each other's prices even though consumers can is sometimes criticised: why should consumers be better informed than firms? More specifically, the argument goes, a firm could send around spies that enquire about prices from competitors, pretending to be potential customers. This criticism is a result of taking the models too literally: in real-world cartels price discounts are not offered uniformly to all customers (in which case there would be indeed no secrecy), but rather to selected customers (preferably ones with large orders). In the model, in contrast, we abstract from this discrimination of customers and instead assume that prices are uniform. Thus, the implicit assumption here is that firms are able to verify that they face a genuine customer before offering any discount from the list price.

As for the assumption that both firms' demand levels are observable, suppose, for a moment, that a firm could observe only its own demand but not the rival's (as Stigler assumed). Price wars might then be initiated by either of the firms on grounds of - not verifiable - unusually low sales due to cheating on part of the other firm. But given that own demand is private information, a firm would *ex post*, after observing a very low demand, never have an incentive to lead the firms into a costly price war, even if it were certain that the other firm had indeed deviated.[13] As a result, collusion could never be sustained in an ordinary Nash equilibrium. An analysis of such a game with imperfect private information would therefore require the use of a different equilibrium notion, and would lead

[13] Cf. the discussion in Fudenberg and Levine (1991).

to severe technical complications.[14] Such problems are circumvented if we assume that the occurence of price wars is conditioned on realisations of a *public* signal, in our case, the vector of quantities.

Moreover, an economic argument in favour of this assumption is that both for tacitly colluding firms and for organized cartels, shipped quantities are likely to be better observable than the accompanying monetary flows (cf. Ulen[1978, p.128], who describes the efforts that the Joint Executive Committee undertook to monitor both shipments and billing, although apparently cheating remained a possibility).

4.3.2 Structure of collusive strategies

The collusive strategies used in the imperfect-monitoring framework introduced above are assumed to belong to a class of Green-Porter-(1984)-type trigger strategies, modified and generalized for a price-setting game with differentiated goods and imperfectly correlated demand shocks.

Roughly, these strategies are characterized by: a cartel price, a set of quantity vectors the occurence of which triggers a price war (this is the set-valued counterpart of the trigger price in Green-Porter, and I will call it the trigger set), and the length of such punishment periods during which the static Bertrand equilibrium is played.

Both for the sake of generality and because it will be relevant later on, we allow that the support of the demand vectors $\mathbf{q}$ that can be observed depends on the price vector $\mathbf{p}$. This is by assumption ruled out in Green-Porter and related works, but arises for example if the demand shocks have compact support (cf. the next subsection).

If this is the case, the space of all possible quantity vectors is partitioned into two subsets, the set of vectors that occur with positive density if both firms

[14] On the analysis of games with imperfect private information, see Fudenberg and Levine (1991) and Lehrer (1992). An oligopoly model with private but perfect information (arising due to localized competition) is studied by Verboven (1995).

adhere to the cartel price, and the set of vectors that can only be observed if (at least) one firm deviates.

A complication now arises. If the trigger set contains vectors from both these subsets, it seems obvious to allow firms to distinguish between (and punish differently) those instances in which it is certain that a firm has deviated, and those for which it is only likely that a deviation has occured.

More formally, denote by Σ the set of all possible quantity vectors (for any price vectors), and by $S(\mathbf{p^c})$, where $\mathbf{p^c} = (p^c, p^c)$, the support of $\mathbf{q}(\mathbf{u})$ in case both firms adhere to the cartel price p^c (thus, models with a price-independent support correspond to the special case $\Sigma = S(\mathbf{p^c})$). According to the above considerations, a collusive strategy would be characterized by the quintuple (p^c, A, T^A, B, T^B). Here, $A \subset \Sigma - S(\mathbf{p^c})$ is a set of demand vectors that cannot occur if both firms charge p^c and that trigger a punishment phase of length T^A, and $B \subset S(\mathbf{p^c})$ contains vectors that can occur without cheating, and which lead to a price war of length T^B.

While in principle strategies of this class could have a very complicated structure, we shall confine the attention to symmetric strategies which maximise the firms' discounted payoffs. Here, the requirement of symmetry refers to symmetry of the sets A and B in the sense that if a vector (q_1, q_2) is contained in a set, then the vector (q_2, q_1) is contained in it as well.[15] Moreover, I assume that on the support $S(\mathbf{p})$, the distribution of $\mathbf{q}$ is characterized by a continuous density function $f(\mathbf{q}, \mathbf{p})$. Given this assumption, we can without loss of generality assume that B is compact. It turns out that the strategies which are optimal within this class have a very simple structure, which is simplified even more if additional assumptions about the demand shocks are introduced.

First, it is clear that it is optimal to punish the occurence of *every* vector that can only be observed if a firm has deviated. Moreover, since in equilibrium

[15] While it is straightforward that for A and B to be optimal, the no-deviation constraints must be binding for both players at the same time (because otherwise these sets could be decreased with a gain in payoff), this does not by itself imply symmetry of the sets.

such events are never observed, it is optimal to punish them maximally. Thereby, maximal compliance is achieved without involving any cost for the cartel. Thus, formally, we have $A = \Sigma - S(\mathbf{p^c})$ and $T^A = \infty$.

A far more difficult question is how to characterize the optimal set B of vectors that occur with positive density in equilibrium, but still lead to a price war. A (partial) answer to this is given in the following proposition.

To state the result, some notation is needed. For some vector $\mathbf{q} = (q_1, q_2)$ denote by $\bar{\mathbf{q}}$ its symmetric counterpart (q_2, q_1). Moreover, define

$$\eta(\mathbf{q}) = \frac{\frac{\partial f}{\partial p_1}(\mathbf{q}, \mathbf{p^c}) + \frac{\partial f}{\partial p_1}(\bar{\mathbf{q}}, \mathbf{p^c})}{f(\mathbf{q}, \mathbf{p^c}) + f(\bar{\mathbf{q}}, \mathbf{p^c})},$$

which is well-defined for any $\mathbf{q}$ in the interior of $S(\mathbf{p^c})$. Then we have

Proposition 4.1 *Let $\eta_0 = \max\{\eta(\mathbf{q}) | \mathbf{q} \in B\}$, and assume that the incentive constraint with respect to marginal deviations from p^c is binding.*[16]

1. Then B is optimal only if it contains all *vectors $\mathbf{q}$ for which $\eta(\mathbf{q}) < \eta_0$ (ignoring any zero-measure subsets)*

2. Moreover, given 1., B must either a) contain all vectors $\mathbf{q}$ for which $\eta(\mathbf{q}) = \eta_0$ as well, or b) inclusion of these vectors in B is not payoff-relevant (i.e. in this case starting with any B that satisfies 1., adding or excluding any vectors $\mathbf{q}$ for which $\eta(\mathbf{q}) = \eta_0$ leads to a payoff-equivalent strategy at the same cartel price and a possibly different punishment duration.)

For all proofs in this chapter, see the Appendix (Section 4.8).

Proposition 4.1 states the important result that the two-dimensional set B, if optimal, can be characterized simply by some scalar η_0. More precisely, an optimal set B contains all pairs of vectors $(\mathbf{q}, \bar{\mathbf{q}})$ for which the relative change in likelihood if a firm marginally deviates from p^c, $(\partial f / \partial p_i)/f$, is less than η_0 (or in

[16] This additional assumption relates to the fact that the incentive constraint for marginal deviations need not be binding for a collusive strategy to be optimal, if (1) the first-best solution (sustaining the monopoly price, and no price wars in equilibrium) can already be sustained, or if (2) the binding incentive constraint relates to large price cuts. This will become clearer in Section 4.3.4 below.

absolute terms exceeds η_0, since for any optimal strategy B η_0 will be negative[17]). In other words, we have the intuitively plausible result that an optimal trigger set B contains vectors that, other things equal, have a small density and/or for which the (negative) derivative with respect to deviations from p^c is large. The result is not only quite powerful in its characterization of optimal trigger sets, but also general in the sense that it refers to the class of trigger strategies introduced here, and does not rest on any additional distributional assumptions.

While the class of strategies considered here is not globally optimal in the sense of Abreu, Pearce and Stacchetti (1986), it is more general than the strategies of Green and Porter (1984) not only with respect to multidimensionality of the demand shocks and the observable signals: while in Green-Porter the trigger set (of prices) is by assumption an appropriately chosen tail of the distribution, here, according to Proposition 4.1, a similar structure of the trigger set emerges (of course depending on the specific distribution) as the necessary characteristic of any optimal strategy.

As will be seen in Section 4.3.4 below, for the analysis of collusion we need to derive explicit collusive strategies. With this requirement, there are two reasons why it is not convenient to consider strategies that are globally optimal in the sense of Abreu, Pearce and Stacchetti (1986, 1990). First, their assumption of a price-independent support of observable signals makes it very difficult to construct examples with analytically tractable distributions. On the other hand, relaxing this assumption alters the problem of finding an optimal strategy in a fundamental way. Second, with a two-dimensional set of observable signals, it is not even in principle clear how to compute the optimal sets that trigger switches between the collusive and the punishment mode. This illustrates the importance of Proposition 4.1.

[17] Otherwise B would contain a subset such that a deviation by a firm would lead to a *decrease* of the probability of such an event, which cannot be optimal.

4.3.3 Demand shocks with compact support

For the demand shock $\mathbf{u}$ introduced in 3.1, I will in this chapter consider a specific family of distributions which has two salient features. First, the correlation between u_1 and u_2 is an increasing function of σ, which corresponds to our key assumption motivated in Section 4.2: the more the products are differentiated, the less the demand shocks are correlated. Second, the variables u_i are assumed to have compact support, which has important consequences for the structure of optimal collusive strategies.

More specifically, we assume that u_1 and u_2 are convex combinations of two i.i.d. variables v_1 and v_2:

$$
\begin{aligned}
u_1 &= \frac{s(1-\sigma)+\sigma}{s(1-\sigma)+2\sigma} v_1 + \frac{\sigma}{s(1-\sigma)+2\sigma} v_2 \quad \text{and} \\
u_2 &= \frac{\sigma}{s(1-\sigma)+2\sigma} v_1 + \frac{s(1-\sigma)+\sigma}{s(1-\sigma)+2\sigma} v_2,
\end{aligned} \tag{4.3}
$$

where $s \geq 0$ is a shift parameter, and v_1 and v_2 are independent and uniformly distributed on an interval $[\mu - d, \mu + d]$ around some mean μ. The parameter $d \geq 0$ determines the variance of the variables v_i and u_i and therefore measures the degree of uncertainty. This will enable us later on to compare models with uncertainty to their deterministic counterparts simply by taking the limit $d \to 0$.

For a more compact notation, I will occasionally write u_1 and u_2 as $u_1 = ((r-\sigma)/r)v_1 + (\sigma/r)v_2$ and $u_2 = (\sigma/r)v_1 + ((r-\sigma)/r)v_2$, where $r = s(1-\sigma)+2\sigma$.

From (4.3), the correlation ρ_u between u_1 and u_2 is easily computed as

$$\rho_u = \frac{2\sigma(s+\sigma-s\sigma)}{s(1-\sigma)[s(1-\sigma)+2\sigma]+2\sigma^2},$$

which is plotted in Figure 4.2 as a function of σ for different values of s. The distributions characterized by the family (4.3) have the following properties:

1. For any positive value of s, the correlation ρ_u is an increasing function of σ, where $\rho_u = 1$ for $\sigma = 1$ and $\rho_u = 0$ for $\sigma = 0$.

2. For any positive value of s, the variance of u_i is a decreasing function of σ, where $\text{Var}(u_i) = 1/2\text{Var}(v_i) = 1/3d^3$ for $\sigma = 1$ and $\text{Var}(u_i) = \text{Var}(v_i) = 2/3d^3$ for $\sigma = 0$.

Figure 4.2: Correlation ρ between u_1 and u_2.

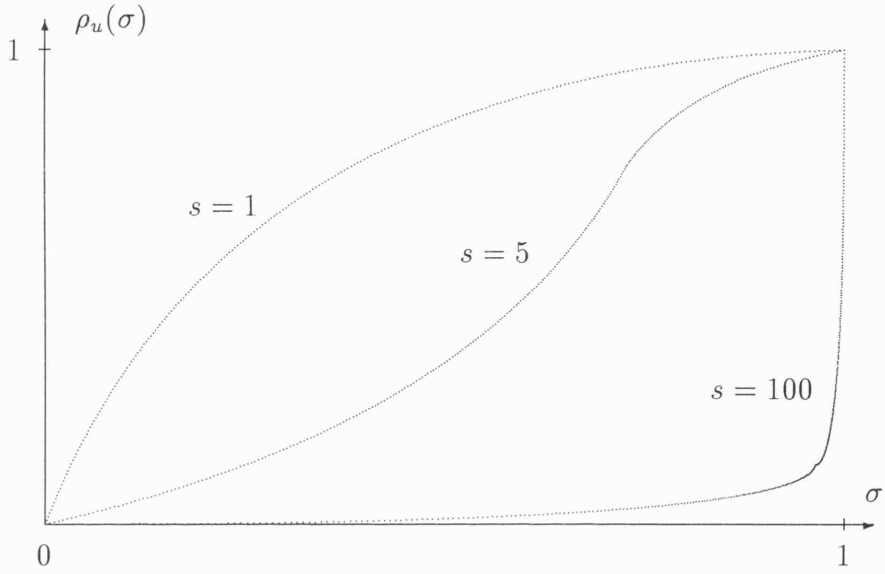

3. The shift parameter s determines how quickly the correlation drops to a low level as goods become more differentiated. It hence plays a similar role to the heterogeneity parameter μ in the model of the previous section, as can be seen by comparing Figures 4.1 and 4.2. In particular, for higher values of s, the schedule $\rho_u(\sigma)$ has a shape very similar to the shape of $\rho_u(\sigma)$ in Figure 4.1 for small values of μ, i.e. little taste heterogeneity.

As for the limit cases, if $s = 0$, then u_1 and u_2 are identical, hence perfectly correlated for all values of σ, and for $s \to \infty$, $u_i = v_i$ $(i = 1, 2)$, i.e. u_1 and u_2 are independent for all values of σ. This latter limit case is important in order to determine in how far the results on the stability of collusion obtained further below depend on the assumption that ρ_u depends on σ.

The second feature of $\mathbf{u}$, the uniform distribution of the underlying variables v_i over a compact interval, is important because a simple application of Proposition 4.1 immediately implies that with this distribution, the optimal B (in the notation of 3.2) is empty, i.e. that only demand vectors *outside* the support $S(\mathbf{p^c})$ should lead to a price war:

Proposition 4.2 *Given the demand shocks (4.3) and their induced distribution* *of* $\mathbf{q}(\mathbf{u}, \mathbf{p^c})$*, it is optimal only to punish the occurence of demand vectors* $\mathbf{q}$ *not contained in* $S(\mathbf{p^c})$.

Hence it is optimal to punish only the occurence of those demand vectors which can only be observed if a firm deviates. The reason for this is of course the rectangular shape of the distribution of v_1 and v_2: given that there is a one-to-one mapping from $\mathbf{v} = (v_1, v_2)$ into $\mathbf{q}$ for any price vector, it follows that all vectors in the interior of $S(\mathbf{p^c})$ have the same density and also the same relative change in likelihood in case of a deviation. Then from Proposition 1 it follows that either all of these vectors should lead to a price war, which clearly is not optimal, or none of them. The result is intuitively clear: vectors $\mathbf{q} \in S(\mathbf{p^c})$ have positive density, but this density hardly changes with marginal changes of the price. Hence, inclusion of such vectors in the trigger set would merely enter negatively into the payoff function without any gain in terms of a more effective deterrence of deviations.

It immediately follows from Proposition 4.2 that with optimal strategies, in this setting, price wars never occur in equilibrium. Related to this, the fact that a price war can only occur if a firm deviates implies that the price war probability is equal to the probability of detection of a deviation (whereas in Green-Porter, there are always type I and type II errors in the inference of deviant behaviour). In order to characterize optimal collusive strategies, it remains to specify the cartel price p^c. We will therefore ask in subsequent sections: what is the maximal price that can be sustained, depending on the degree of product differentiation?

Now the analysis may look rather like that of the standard deterministic models: find the maximal sustainable cartel price given that firms use a simple grim trigger strategy. But this is still an imperfect-information environment in a fundamental sense. First, the fact that a price war can only occur if a firm in fact deviates is not a simplifying assumption but is derived as part of an optimal strategy, given the class of analytically convenient distributions used here. Second, and more importantly, while a price war can only occur if a firm deviates, this of course does not imply that any deviation will indeed be detected.

Rather, marginal increases of the price war/detection probability (from zero) will have to suffice to deter marginal deviations.

4.3.4 Payoff function and incentive constraints

In this section, we derive the two central incentive constraints which determine the sustainability of a collusive strategy. These constraints will subsequently be used in the models of Sections 4.4 and 4.5 in order to analyse the stability of collusion as a function of product differentiation.

We have seen in Section 4.3.3 that an optimal strategy is simply characterized by the cartel price p^c. A price war occurs if any vector $\mathbf{q} \in \Sigma - S(\mathbf{p^c})$ is observed, which leads to maximal punishment, i.e. breakdown of the cartel. Assuming that firm 2 adheres to the cartel price p^c, denote the probability of a price war by $\alpha(p_1, p^c) = \text{Prob}\{\mathbf{q}(p_1, p^c) \in \Sigma - S(\mathbf{p^c})\}$. While $\alpha = 0$ if both firms adhere to p^c, it also follows from the assumptions on the distribution of the demand shocks that for p_1 sufficiently low, we can have $\alpha = 1$, viz. if a deviation leads to a demand vector outside $S(\mathbf{p^c})$ with certainty.

Denoting firm 1's per-period profit by $\pi_1(p_1, p^c)$, firm 1's expected discounted payoff from an infinitely repeated game is characterized by the Bellman equation

$$v_1(p_1, s^c) = \pi_1(p_1, s^c) + [1 - \alpha(p_1, s^c)]\delta v_1(p_1, s^c) + \alpha(p_1, s^c)\frac{\delta}{1-\delta}\pi^b,$$

where π^b is the per-period profit in the static Bertrand equilibrium. This equation can be be solved explicitly for v_1:

$$v_1 = \frac{\pi^b}{1-\delta} + \frac{\pi_1 - \pi^b}{1 - \delta + \delta\alpha}, \tag{4.4}$$

where the arguments of v_1, π_1 and α have been omitted (cf. Green and Porter 1984).

If firms adhere to the cartel price, then $\alpha = 0$, and the collusive payoff is simply $\pi^c/(1-\delta)$, with $\pi^c = \pi_1(\mathbf{p^c})$. Sustainability of the cartel price p^c requires that v_1 as given by (4.4) has a global maximum at $p_1 = p^c$, or

$$\frac{\pi^c - \pi^b}{1-\delta} \geq \frac{\pi_1(p_1, p^c) - \pi^b}{1 - \delta + \delta\alpha(p_1, p^c)} \quad \forall\, p_1,$$

which leads to the general incentive constraint

$$\frac{\pi_1(p_1,p^c)-\pi^b}{\pi^c-\pi^b}\leq\frac{\delta}{1-\delta}\alpha(p_1,p^c)\quad\forall p_1. \tag{4.5}$$

This incentive constraint includes as special cases two more specific and simpler incentive constraints, which much of the subsequent analysis will be concerned with.

The first states that marginal deviations must not be profitable. Differentiation of (4.5) with respect to p_1 at the price p^c (at which both sides of (4.5) vanish) leads to the constraint

$$\theta^m:=\frac{\partial\pi_1(\mathbf{p^c})/\partial p_1}{(\pi^c-\pi^b)\partial\alpha(\mathbf{p^c})/\partial p_1}\leq\frac{\delta}{1-\delta}. \tag{4.6}$$

This constraint, of course, is familiar from Green and Porter (1984) and related works: the marginal gain of deviating from p^c must be counterbalanced by an expected loss in payoff due to an increase of the price war (= detection) probability α in order to deter deviations.

The second constraint relates to the case where a deviating firm sets p_1 at such a low level that the deviation is detected with certainty, i.e. if any realised demand vector $q(p_1,p^c)$ lies outside $S(\mathbf{p^c})$, which can occur with distributions with compact support. Given that $\alpha=1$, i.e. that deviation leads to the breakdown of the cartel for sure, a deviating firm would then want to maximise its current-period profit. That is, it would set the price $p^d(p^c)$ which maximises $\pi(p_1,p^c)$. If we denote the resulting defection profit by π^d, the incentive constraint now becomes

$$\theta^d:=\frac{\pi^d-\pi^c}{\pi^c-\pi^d}\leq\frac{\delta}{1-\delta}. \tag{4.7}$$

This is the incentive constraint familiar from the theoretical literature discussed in the Introduction: given that a deviation leads to retaliation with certainty, the incentive constraint states that the ratio of the gain from deviating to the benefit of colluding must not exceed a critical level which depends on the discount factor.

Now, it is important to what follows that neither of the constraints (4.6) and (4.7) implies the other. This is due to the fact that the payoff function $v_i(p_i)$

Figure 4.3: Payoff function $v_1(p_1, p^c)$

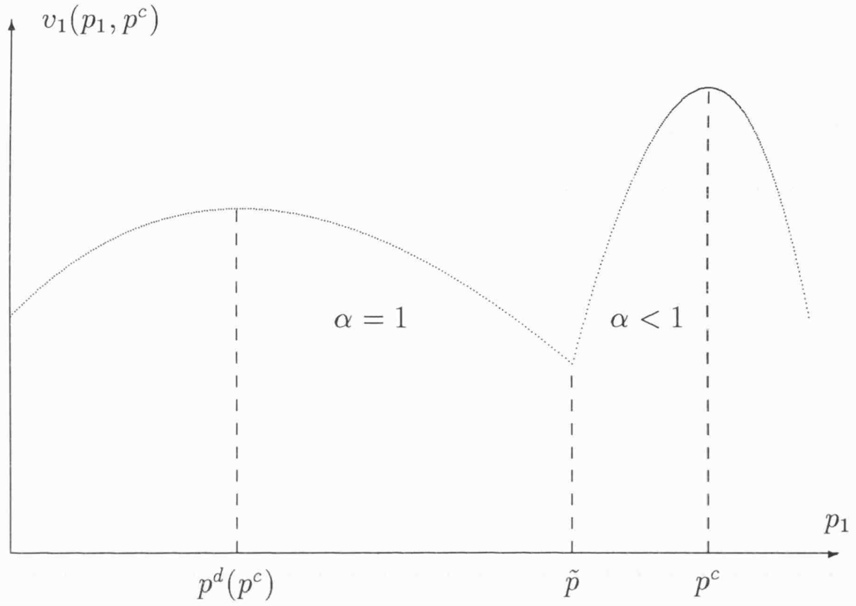

is necessarily nonconcave where α is at or close to one, as depicted in Figure 4.3: To see why this must be the case, consider two price regions, prices not too far below the cartel price p^c for which a deviation is not necessarily detected ($\alpha < 1$), and lower prices for which breakdown follows with certainty ($\alpha = 1$). If (4.6), the incentive constraint for marginal deviations, is satisfied, then v might be increasing in p_1 as long as $\alpha < 1$. That is, the loss in payoff due to an increase of α increasingly outweighs the gain from deviating.

But this trade-off disappears once the level $\alpha = 1$ is reached: while v will exhibit a local minimum at the price $\tilde{p}$ depicted in Figure 4.3, it might be profitable to decrease p_1 further, i.e. a firm would want to maximise its current-period profit, given that in this range of prices breakdown of the cartel follows anyway. Hence, v has another local maximum at $p^d(p^c)$, as depicted in the figure.

It is important to see that the nonconcavity of v is not a consequence of the particular distributions considered here, but a very general phenomenon which is also present in models like Green and Porter (1984), where α can never take the values 0 or 1: while it is true that, according to Porter's (1983) Proposition 2.3, convexity of α in p_1 and concavity of π_1 imply that v is increasing (as in Figure

4.3 in the ($\alpha < 1$)-region), α cannot be convex over the entire relevant range of prices if its density is concentrated in a narrower band of prices around the cartel price. And in this case, the same situation as discussed above might arise: while a marginal deviation may not be profitable, a firm might want to deviate to a very low price, even if then retaliation follows almost with certainty. Hence, while the kink depicted in Figure 4.3 is a consequence of the distributions used here, the nonconcavity of v is general.

From the above discussion it follows that provided (4.6) implies that (4.5) is satisfied for any prices for which $\alpha < 1$, the two conditions (4.6) and (4.7) are *sufficient* for the general condition (4.5), i.e. for sustainability of p^c, which implies a considerable simplification for the analysis. For future reference we state this condition formally:

Condition 4.1 *If for a given price p^c (4.6) holds, then (4.5) is satisfied for all p_1 such that $\alpha(p_1, p^c) < 1$.*

4.4 Example 1: a Hotelling model

In this and the next section, we apply the framework set out in Section 4.3 to two specific examples, in order to see how the sustainability of collusion, as determined by the two central incentive constraints (4.6) and (4.7) derived above, depends on the degree of product differentiation. The first model is a Hotelling model quite similar to the one studied in Section 4.2. In contrast to that model, however, the correlation of demand functions as a function of product differentiation will be assumed rather than derived.

4.4.1 The stage game

As in the model of Section 4.2, two firms are located symmetrically on a line $[0, 1]$, firm 1 at $\sigma/2$ and firm 2 at $1 - \sigma/2$. Consumers are distributed along the line. The utility of a consumer located at z that purchases firm i's product is given by

$$w_i(z) = y + a - p_i - b(z - z_i)^2, \tag{4.8}$$

where y denotes income, a is the utility derived from consuming the most preferred good, b is a parameter for "travel costs", p_i is firm i's price and z_i its location. The parameter a is assumed to be sufficiently large such that the market is fully covered at all relevant prices. This condition will be stated more precisely below.

While individual utility is deterministic here, there are again "macro"shocks affecting the densities of consumers in group 1 in the interval $[0, 1/2)$, and of group 2 in $(1/2, 1]$. These densities u_1 and u_2 are assumed to belong to the family of distributions (4.3), with mean value $\mu = 1$. That is, u_1 and u_2 are convex combinations of two i.i.d. variables which are uniformly distributed over $[1 - d, 1 + d]$, and the correlation of u_1 and u_2 is an increasing function of σ.

The assumption of correlated density shocks is itself arbitrary, but it gives rise to correlated demand functions in the same way as the model of Section 4.2. Thus it is the model of Section 4.2 that provides the justification for this assumption.

From (4.8), one can calculate the location of the marginal consumer between firms 1 and 2, and from that the firms' demand functions. These demand functions are functions of the price difference $k := p_1 - p_2$, and in what follows we will assume without loss of generality $k \leq 0$, i.e. it is always firm 1, if any, that undercuts the other firm. Then we have

$$q_1 = \frac{1}{2}u_1 - \frac{k}{2t}u_2 \qquad \text{and} \qquad q_2 = \frac{t + k}{2t}u_2, \tag{4.9}$$

where $t := b(1 - \sigma)$.

Firm 1's expected demand is $\mathrm{E}q_1(k) = (t - k)/(2t)$, and if firm 2 charges a price p, firm 1's expected profits are $\pi_1(k, p) = (k + p)(t - k)/(2t)$. Solving the first-order condition $\partial\pi_1(0, p)/\partial k = 0$ for p then yields the Bertrand equilibrium price $p^b = t$.

Under the assumption that full market coverage is always optimal, the joint profit maximising price (hereafter called monopoly price and labelled by 'm') is determined by setting the utility of the consumer furthest away from a firm to zero. For firm 1, this consumer is $z = 1/2$ if $\sigma \leq 1/2$, and $z = 0$ for $\sigma \geq 1/2$, i.e. if firms are located more towards the centre (cf. Chang 1991). We then get

$$p^m = a - b\left(\frac{1-\sigma}{2}\right)^2 \text{ for } \sigma \in [0, 1/2], \quad \text{and} \quad p^m = a - b\left(\frac{\sigma}{2}\right)^2 \text{ for } \sigma \in [1/2, 1].$$

For the first-order approach of the calculation of the Bertrand price above to be correct, we need $p^m \geq p^b$ for all values of σ. A sufficient condition is that $p^m \geq p^b$ at $\sigma = 0$, which leads to the parameter restriction $a \geq (5/4)b$. Finally, starting at p^m, an increase of price by $\mathrm{d}p$ leads to a gain of $\mathrm{d}p/2$ for each firm (since the expected consumer density is one) and to a loss of $p^m \mathrm{d}p$. Hence, that full market coverage indeed be optimal leads to a second parameter restriction $4a - b \geq 2$.

4.4.2 Collusion: discrete deviations

In this section, we analyse the requirement (4.7) that deviations from a collusive price by a large amount, i.e. such that detection is certain, be unprofitable. To evaluate θ^d, the l.h.s. of (4.7), we need to know the expected per-period profits in the Bertrand equilibrium and with collusion. These are simply given by $\pi^b = p^b/2$ and $\pi^c = p^c/2$, respectively. The optimal deviation price $p^d(p^c)$ (cf. Section 4.3.4), here expressed in terms of the price difference k^d, is in a first step obtained as the solution to the first-order condition $\partial\pi_1(k, p^c)/\partial k = 0$, which is $k_+^d = -(p^c - t)/2$. On the other hand, by setting $k = k_0^d := -t$, firm 1 takes over the whole market (cf. (4.9)). Hence, the optimal deviation is given by k_+^d, provided that $k_+^d \geq k_0^d$ or equivalently $p^c \leq 3t$, and otherwise by k_0^d. In other words, for sufficiently homogeneous goods, i.e. higher values of σ (and hence small t), an optimally deviating firm will set p^d so as to take over the whole market, whereas for more differentiated goods, optimal deviation will still leave positive demand for the other firm.

Denote by θ_+^d and θ_0^d the relevant expressions for θ^d, given that firm 1 deviates by k_+^d and k_0^d, respectively. Then we have for the case $p^c \geq 3t$, $k^d = -t$, $\pi^d = p^c - t$ and then $\theta_0^d = 1 - t/(p^c - t)$, and if $p^c \leq 3t$, $k^d = -(p^c - t)/2$, $\pi^d = (p^c + t)^2/(8t)$, and $\theta_+^d = (p^c - t)/(4t)$.

Consider a fixed cartel price p^c. With $t = b(1-\sigma)$, there exists a critical $\sigma_0 \geq 0$ such that deviation to k_0^d is optimal (and hence θ_0^d relevant) for all $\sigma \geq \sigma_0$, and deviation to k_+^d for all $\sigma < \sigma_0$. This σ_0 is strictly positive if and only if $p^c < 3b$. This defines the schedule $\theta^d(\sigma)$, and it is easy to see that both θ_+^d and θ_0^d and hence θ^d are increasing functions of σ.

Now consider the case $p^c = p^m(\sigma)$. Since p^m is decreasing in σ for $\sigma > 1/2$, a nonmonotonicity could possibly arise with respect to $\theta(\sigma)$. But this is not the case: with $p^c = p^m(\sigma)$, there still exists a critical σ_0 with the above properties, and since by assumption $a > b$, both θ_+^d and θ_0^d are still increasing functions of σ.

According to the standard interpretation of critical discount factors or rates, the fact that θ^d is increasing in σ means that, as far as the constraint (4.7) is concerned, a given cartel price p^c is easier to sustain for differentiated goods than for more homogeneous goods, in the sense that the incentive constraint $\theta \leq \delta/(1-\delta)$ is less restrictive. The same relationship has been derived for several other quite different models in the theoretical literature discussed in the introduction of this chapter. It is this relationship that has led many theorists to question the validity of the traditional view. Here, it is not entirely clear whether the recurrence of this picture is a merely a coincidence or has some more fundamental origin. All we can say is that while both the collusive benefit and the gain from deviating are usually larger for homogeneous goods than for differentiated goods, in a variety of models the gain from deviating increases *proportionally faster* than the collusive benefit as goods become more homogeneous, thereby leading to a decrease of cartel stability, contrary to the traditional view.

The maximum of θ^d is attained at $\sigma = 1$, where $\theta = 1$. It follows that for $\delta > 1/2$, the constraint (4.7) is satisfied for all values of σ. In other words, the incentive constraint for discrete deviations is relevant only if the discount factor is relatively low.

4.4.3 Collusion: marginal deviations

Now we turn to the analysis of the incentive constraint (4.6) which relates to marginal deviations from the cartel price p^c. In order to evaluate θ^m, the expression on the l.h.s. of (4.6), we need to know the detection probability α, which is stated in the following lemma as a function of k and p^c, the price set by firm 2:

Lemma 4.1 *The probability* $\alpha(k,p^c) = Prob\{\mathbf{q}(k,p^c) \in \Sigma - S(\mathbf{p^c})\}$ *is given by*

$$\begin{aligned}\alpha(k,p^c) &= -\frac{k[(r+2\sigma)t+(r+\sigma)k]}{2(t+k)[(r-2\sigma)t-k\sigma]} \\ &- \frac{kr(2t+k)[kr+3dkr-4dk\sigma+4dt(r-2\sigma)]}{8d^2(k+t)[(r-2\sigma)t+k(r-\sigma)][(r-2\sigma)t-k\sigma]}(1-d).\end{aligned}$$

From Lemma 4.1, the marginal increase of the probability of detection at $k=0$ is

$$\frac{\partial\alpha(0,p^c)}{\partial k} = -\frac{s(2-d)(1-\sigma)+4\sigma}{2dsb(1-\sigma)^2}. \tag{4.10}$$

Moreover, we have $\pi^c - \pi^b = (p^c - t)/2$ and $\partial\pi_1/\partial k = -(p^c-t)/(2t)$, and the ratio of these two expressions is simply $-1/t$, i.e. does not depend on p^c. Therefore, the left-hand side of the incentive constraint (4.6) for marginal deviations is

$$\theta^m = \frac{2ds(1-\sigma)}{s(2-d)(1-\sigma)+4\sigma}. \tag{4.11}$$

It is easy to see that θ^m is a decreasing function of σ. Hence, as far as marginal deviations from the cartel price are concerned, collusion is *less* sustainable for differentiated goods than for homogeneous goods, in accordance with the traditional view.

To understand this relationship, consider the case of homogeneous goods. Here, the marginal gain from deviating may be rather large, but this is outweighed by an even larger (in absolute terms) increase of the detection probability. The reason for this is that, according to our assumption that the demand functions are highly correlated for equal prices, a deviation is relatively likely to lead to realisations of the demand vector outside the permissable region $S(\mathbf{p^c})$. This relative undesirability of marginal price cuts for large σ is reflected in a low value of θ^m.

Conversely, for more differentiated goods, the deterrent force of an increase of α vis-à-vis the increase of current-period profits is weaker, since a lower correlation of the demands corresponds to a larger set of demand vectors that occur with positive density. A price cut, therefore, is less likely to result in a realised demand vector outside this region.

The effect of a decrease in the correlation due to differentiation is reinforced by the accompanying increase in the variance of the demand shocks. In the Section 4.5.3, I will discuss how the relative importance of the correlation effect and the change in the variance can be assessed.

The price and cross-price elasticities of demand, which of course vary with σ as well, play a double role, in such a way that in this model they do not affect the incentive constraint (4.6) at all: with a high elasticity in the case of homogeneous goods, marginal deviations may be tempting. But precisely the ability for a firm to capture a large share of the market by cutting the price only by a small amount also implies that such a deviation is likely to be detected. In this particular model, now, these two effects exactly cancel each other. This can be seen by considering the limit case $s \to \infty$, which corresponds to the case that the density shocks u_1 and u_2 are uncorrelated (cf. Section 4.3.1). Here, we have $\theta^m = 2d/(2-d)$. Since this expression does not depend on σ, this verifies that a negative relationship betweeen θ^m and σ depends on the assumed correlation of the demand shocks, but not on the elasticity of demand.

The comparative-statics properties of θ^m with respect to s and d are as follows: an increase in s means that the correlation of the demand shocks decreases faster as products become more differentiated. This is reflected in an upward shift of $\theta^m(\sigma)$, i.e. a decrease of cartel stability. An increase an d has a similar effect, but for a different reason: while d does not affect the correlation of the demand shocks, it determines their variance. A more noisy environment is reflected in a decrease of α for any given σ and thus renders marginal deviations more profitable, which is reflected in an upward shift of $\theta(\sigma)$.

It may perhaps be surprising that θ^m and hence the incentive constraint for marginal deviations does not depend on the cartel price (since p^c drops out in the ratio $(\partial\pi_1/\partial k)/(\pi^c - \pi^b)$). This can be shown to be a general property of any model in which demand is an affine function of the difference of prices and is therefore a consequence of the Hotelling framework. The important implication is that the incentive constraint (4.6) here refers to sustainability of collusion *as such* and not only with reference to a particular collusive price. This issue will be discussed again further below.

Before we can go on to examine both incentive constraints (4.7) and (4.6) together, we have to see whether Condition 4.1 is satisfied, i.e. the condition that if marginal price cuts are not profitable, then no other price cuts for which the resulting detection probability is less than one, be profitable. Here we face the problem that $\alpha(k, p^c)$ can be either convex or concave in k. In particular, it turns out that $\alpha(k, p^c)$ is always concave at $k = 0$. But, with both π_1 and α concave in k at 0, Condition 4.1 can only be satisfied if the profit function is "more concave" at 0 than α, i.e. if $\partial^2\pi(k, p^c)/\partial k^2 \leq \partial^2\alpha(k, p^c)/\partial k^2$. Otherwise, if (4.6) were binding, then a small discrete deviation would be profitable. That is, satisfaction of (4.6) would not imply satisfaction of the general constraint (4.5).

The following result shows that for this requirement to be met, p^c must not exceed a certain upper bound which depends negatively on s:

Lemma 4.2 *Assume (4.6) is binding, and let $\beta = \delta/(1-\delta)$. Then the condition $\partial^2\pi(k, p^c)/\partial k^2 \leq \partial^2\alpha(k, p^c)/\partial k^2$ at $k = 0$ is equivalent to*

$$p^c \leq p^c_{max} := \frac{4b\beta(2+\beta)^2(1+d^2)}{(4 - 12\beta + \beta^2 + 4d^2 + 4\beta d^2 + \beta^2 d^2)(4\beta - 2\beta s + 2ds + \beta ds)}. \tag{4.12}$$

A sufficient condition for this is

$$p^c \leq \frac{4b\beta(2+\beta)^2}{(2-\beta)^2(4\beta + 2s - \beta s)}, \tag{4.13}$$

and for $\delta \geq 1/2$, a sufficient condition for (4.13) is $p^c \leq 36b/(4+s)$.

Moreover, even where the condition of Lemma 4.2 is satisfied, Condition 4.1 could still be violated, depending on the behaviour of $\partial^3\pi(k, p^c)/\partial k^3$ and $\partial^3\alpha(k, p^c)/\partial k^3$

Figure 4.4: Left-hand sides of the incentive constraints (3.6) and (3.7)

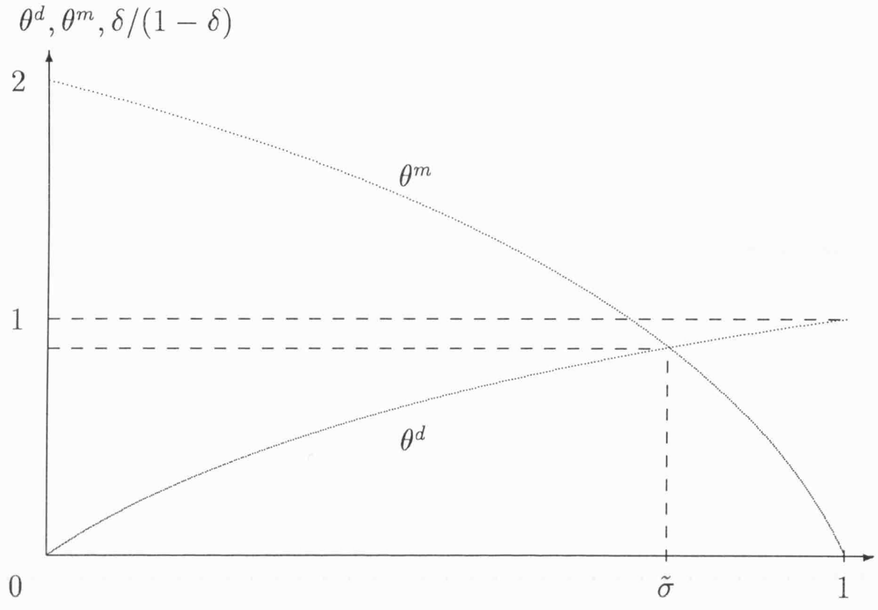

for $k < 0$. Here, we have $\partial^3\pi(k,p^c)/\partial k^3 = 0$, while numerical experiments indicate that $\partial^3\alpha(k,p^c)/\partial k^3 < 0$ for $k < 0$, provided (4.12) holds. This means, while α may be concave over some range of k, the l.h.s. and r.h.s. of (4.5) never intersect again if both (4.6) and (4.12) hold. Therefore, I will in the following assume that Condition 4.1 is satisfied, i.e. that we need not worry further about small discrete price cuts.

4.4.4 Both incentive constraints combined

Having verified that the two constraints (4.6) and (4.7) are sufficient for the sustainability of a collusive price p^c, we can now analyse the stability of collusion by bringing together both constraints. The left-hand sides of the incentive constraints (4.6) and (4.7), θ^d and θ^m, are depicted in Figure 4.4 as functions of σ and for the case of maximal uncertainty, i.e. $d = 1$. Recall that the price p^c is sustainable if $\delta/(1-\delta)$ exceeds both θ^d and θ^m. Define $\tilde{\sigma}$ as the σ solving $\theta^d(\sigma) = \theta^m(\sigma)$, and define $\delta(\sigma)$ as the solution of the equation $\theta^d(\sigma) = \delta/(1-\delta)$. Then from Figure 4.4, we see that there are four different intervals of values of δ:

First, for very low discount factors ($\delta < \delta(\tilde{\sigma})$), p^c is not sustainable for any degree of product differentiation because at least one of the incentive constraints is violated. Similarly, for high discount factors ($\delta \leq 2d/(2+d)$), any price p^c is sustainable for any degree of product differentiation.

The economically more interesting range of discount factors includes those values where the sustainability of a given cartel price depends on σ. Here we have two different cases: for low values of δ ($\delta \in [\delta(\tilde{\sigma}), 1/2)$) such that discrete deviations are relevant, the positive relationship between product differentiation and cartel stability, familiar from other papers, holds only over a certain interval for σ, whereas for larger degrees of differentiation the incentive constraint for marginal deviations implies a decrease in cartel stability. Thus, the sustainability of p^c depends on σ in a nonmonotonic way: p^c may be sustainable for intermediate levels of differentiation, but neither for homogeneous goods nor for larger degrees of differentiation.

For $\delta \in [1/2, 2d/(2+d))$, the incentive constraint for discrete deviations is always satisfied. Here, monotonicity of θ^m in σ implies a monotonic relationship between the sustainability of p^c and the degree of differentiation, i.e., there exists a critical level of differentiation such that for more homogeneous goods, p^c is sustainable, whereas it is not for more differentiated products. This result, the economic intuition of which has already been discussed in Section 4.4.3, is consistent with the traditional view on the sustainability of collusion as a function of product differentiation, and states the opposite of what is usually found in recent theoretical work.

The intercept of θ^m at $\sigma = 0$ depends on the level of demand uncertainty d. A consequence of this is that the monotonicity region $[1/2, 2d/(2+d)]$ for δ is nonempty only if there is sufficient uncertainty, viz. $d > 2/3$. Otherwise, the zones described above look different: there is a region of nonmonotonicity for $\delta \in [\delta(\tilde{\sigma}), 2d/(2+d)]$, and a region of inverse monotonicity (i.e. in the direction predicted in other papers) for $\delta \in [2d/(2+d), 1/2]$. Hence, in this model, whether the predicted relationship between product differentiation and cartel stability is more in line with the traditional view or more with recent theoretical contributions is determined by the level of demand uncertainty.

4.5 Example 2: a representative-consumer model

We now carry out the same analysis as in the previous section for a quite different oligopoly model, one in which the demand system is derived from a representative consumer's utility function. The comparison of the results obtained for this and the previous model will allow us to assess which predictions seem to be fairly robust, and which ones depend on the details of the particular model under consideration.

4.5.1 The stage game

Consider a representative consumer with the linear-quadratic utility function

$$u(\tilde{q}_1, \tilde{q}_2) = a(\tilde{q}_1 + \tilde{q}_2) - \frac{b}{2}\left(\tilde{q}_1{}^2 + 2\sigma\tilde{q}_1\tilde{q}_2 + \tilde{q}_2{}^2\right),$$

where a and b are positive parameters, and $\sigma \in [0,1]$ is the substitutability parameter. While this specification is so far standard, I assume that there is demand uncertainty resulting from randomness in the utility the representative consumer derives, in the sense that the perceived quantity $\tilde{q}_i$ may not coincide with the actual quantity q_i consumed. More precisely, assume that

$$\tilde{q}_1 = q_1 - \frac{u_1}{b(1+\sigma)} \quad \text{and} \quad \tilde{q}_2 = q_2 - \frac{u_2}{b(1+\sigma)},$$

where u_1 and u_2 are the random shocks of equation (4.3) with zero mean, with all the properties discussed there. Proceeding in standard fashion, maximisation of the representative consumer's utility function leads to a system of indirect demand functions

$$\begin{aligned} p_1 &= a - bq_1 - b\sigma q_2 + \frac{u_1 + \sigma u_2}{1+\sigma} \quad \text{and} \\ p_2 &= a - b\sigma q_1 - bq_2 + \frac{\sigma u_1 + u_2}{1+\sigma}. \end{aligned} \tag{4.14}$$

Inversion of (4.14) yields the direct demand functions

$$
\begin{aligned}
q_1 &= \frac{m}{1-\sigma}\left[(1-\sigma)(a+u_1)-p_1+\sigma p_2\right] =: q_1^n \quad \text{and} \\
q_2 &= \frac{m}{1-\sigma}\left[(1-\sigma)(a+u_2)+\sigma p_1-p_2\right] =: q_2^n
\end{aligned}
\tag{4.15}
$$

with $m = b/(1+\sigma)$, for the case that demand for both goods is positive. If $q_j = 0$ $(j = 1,2)$, then demand for the other good i is

$$
q_i = m\left[(1+\sigma)(a-p_i)+u_i+\sigma u_j\right] =: q_i^m. \tag{4.16}
$$

Hence, for given $\mathbf{u}$, the firms' demand functions exhibit a kink at some price. As for the interpretation of the demand shocks in this context, a shock u_i positively affects the demand for good i because this shock negatively affects the quantity $\tilde{q}_i$ the representative consumer perceives (vis-à-vis the actual quantity q_i), and therefore leads to an increase in actual consumption.

As for the correlation of the demand shocks, however, this framework does not provide a foundation for the correlation of the demand shocks being a function of σ, as an alternative to the theory of Section 4.2. Suppose that, in contrast to our assumption, u_1 and u_2 were independent, and consider a value of σ close to one, i.e. nearly homogeneous products. Here, it turns out that according to (4.14), for equal quantities the prices of the goods must be highly correlated, whereas according to (4.15), for equal prices the quantities demanded are independent!

Looking back at the Hotelling model of Section 4.2 might help to explain how this can be the case. If goods are nearly homogeneous, then prices must be nearly the same for both goods to enjoy positive demand. But, conversely, equality of prices does not imply equality of demand: if the heterogeneity parameter μ in the model of Section 4.2 is zero, then the two firms' demand functions are independent because their market areas are disjoint, even if the goods are nearly homogeneous. It is for this reason that here, the dependence of the correlation of demand shocks must be explicitly introduced by assumption, according to the specification (4.3).

Given the demand system (4.15), firm i's expected profit is given by

$$
\pi_i = (p_i - c)\frac{m}{1-\sigma}\left[(1-\sigma)a - p_i + \sigma p_j\right], \tag{4.17}
$$

where c is a constant marginal cost. Given this, it is straightforward to compute prices and profits for the static Bertrand equilibrium (labelled by 'b') and the joint profit maximising solution (labelled by 'm'):

$$\begin{aligned} p^b &= c + \frac{(1-\sigma)(a-c)}{2-\sigma}, \qquad p^m = \frac{a+c}{2}, \\ \pi^b &= \frac{m(1-\sigma)(a-c)^2}{(2-\sigma)^2}, \qquad \pi^m = m\frac{(a-c)^2}{4}. \end{aligned}$$

Moreover, if both firms set the same price p^c, the resulting expected profit is $\pi^c = m(p^c - c)(a - p^c)$.

We assume that for all prices up to the monopoly price p^m, the firms' demands are given by the system (4.15) as long as both firms charge the same price. In other words, a firm's demand is generically positive unless the other firm charges a lower price. Since the worst case for firm i is a shock $u_i = -d$, from (4.15) we have the condition $m(a - d - p^m) \geq 0$, which implies an upper bound for the amplitude of the demand shocks, viz. $d \leq (a - c)/2$.

4.5.2 Collusion: discrete deviations

As in section 4.4.2, we now analyse the left-hand side of the incentive constraint (4.7) for large deviations from p^c that lead to retaliation, i.e. breakdown of the cartel, with certainty. The first part of this analysis, however, will be concerned with a technical complication that arises due to the specific structure of the system of demand functions as given by (4.15) and (4.16).

As in the Hotelling model of the previous section and other models, in the calculation of optimal deviations two cases must be distinguished: the case where the nondeviating firm still enjoys positive demand, and the case where the deviating firm takes over the whole market.

Here we face a technical problem: for a certain range of prices of the deviant, whether the other firm has zero or positive demand depends on the realisation of the demand shock u. That is, if for instance firm 1 deviates, then for higher values of u_2, the relevant demand for both firms might be given by (4.15), while for lower values, firm 2's demand might be zero, in which case firm 1's demand

Figure 4.5: True expected demand $\mathrm{E}_u q_i(\mathbf{p}, \mathbf{u})$ and its approximation $q_i^{app}(\mathbf{p})$

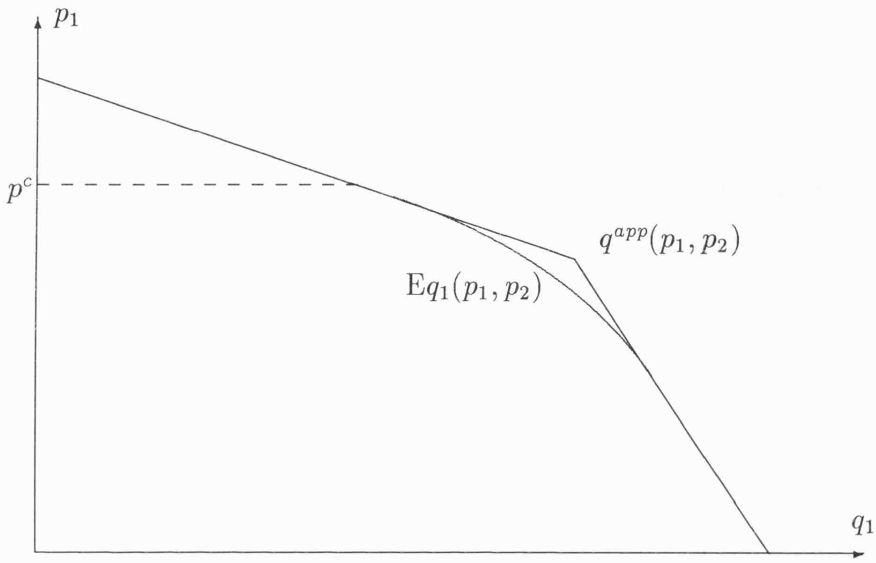

would be given by (4.16). It turns out that the calculation of the correct expected demand for a deviating firm, taking this distinction into account, leads not to a simple linear-kinked curve, but rather to expressions very awkward to use in the subsequent analysis.

In what follows, therefore, we shall simplify the analysis by making use of an upper bound on the precise expected demand, which is simply the maximum of the demand expressions q_i^n and q_i^m, evaluated at $\mathbf{u} = 0$, as the following result states:

Lemma 4.3 *For any price vector p, we have*

$$\begin{aligned} E_u q_i(\mathbf{p}, \mathbf{u}) \leq q_i^{app} &:= \min\{q^n(\mathbf{p}, \mathbf{0}), q^m(\mathbf{p}, \mathbf{0})\} \\ &= \min\left\{\frac{m}{1-\sigma}\left[(1-\sigma)a - p_i + \sigma p_j\right], m(1+\sigma)(a-p_i)\right\}. \end{aligned}$$

This approximation by a linear-kinked curve is relevant only for an intermediate price range, whereas for prices close to p^c and very low prices, the approximation is exact. This relationship between $\mathrm{E}q_i$ and q_i^{app} is depicted in Figure 4.5. I will discuss further below how results derived using q_i^{app} relate to those for the true expected demand, and why this simplification does not affect the results derived below in any significant way.

Using q_i^{app} instead of $\mathrm{E}q_i$, we now analyse the incentive constraint for discrete deviations in precisely the same way as in Section 4.4.2. The prices and profits of a firm that maximises current-period payoff, given that the other adheres to p^c, can be obtained from (4.17) by solving the first-order condition $\partial\pi_1(p_1, p^c)/\partial p_1 = 0$ for p_1:

$$\begin{aligned} p_+^d &= \frac{1}{2}\left[a(1-\sigma)+c+\sigma p^c\right] \quad \text{and} \\ \pi_+^d &= \frac{m}{4(1-\sigma)}\left[a(1-\sigma)+c+\sigma p^c\right]^2. \end{aligned}$$

These expressions are valid as long as for the price vector (p^d, p^c), firm 2's expected demand q_2^{app} is nonnegative. Otherwise, for lower prices p_1, the optimal deviation of firm 1 is given by the equation $q_2^{app}(p_1, p^c) = 0$, which leads to

$$\begin{aligned} p_0^d &= \frac{1}{\sigma}\left[p^c - a(1-\sigma)\right] \quad \text{and} \\ \pi_0^d &= \frac{m}{\sigma}(1+\sigma)(a-p^c)\left[p^c-(1-\sigma)a-\sigma p^c\right]. \end{aligned}$$

The sustainability of a cartel price p^c with respect to large discrete deviations can now be determined by substituting the expressions for π^c, π^b and π^d into the incentive constraint (4.7). For the l.h.s. of this constraint we obtain the expressions

$$\begin{aligned} \theta_+^d &= \frac{(2-\sigma)^2\left[(2-\sigma)p^c-(1-\sigma)a-c\right]}{4(1-\sigma)\left[a+c(1-\sigma)-(2-\sigma)p^c\right]} \quad \text{and} \\ \theta_0^d &= \frac{(a-p^c)(2-\sigma)^2\left[(1-\sigma^2)a-(1+\sigma-\sigma^2)p^c+c\sigma\right]}{\sigma^2\left[a(1-\sigma)-(2-\sigma)p^c+c\right]\left[a-(2-\sigma)p^c+(1-\sigma)c\right]}. \end{aligned} \tag{4.18}$$

As in the previous model, the equation $p_+^d = p_0^d$ defines the value of σ above which the relevant incentive constraint is given by θ_0^d, and below by θ_+^d. This critical value of σ is strictly between 0 and 1, because for $\sigma = 1$, $p_0^d > p_+^d$, and for $\sigma = 0$, a firm can never push the other firm's demand to zero.

As for the resulting schedule $\theta(\sigma)$, we have $\theta = 1$ for both $\sigma = 0$ and $\sigma = 1$. Moreover, in its relevant range, θ_+^d is strictly increasing in σ, whereas θ_0^d is first increasing, but eventually decreasing in σ.

Figure 4.6: Left-hand side of incentive constraint (3.7) for discrete deviations

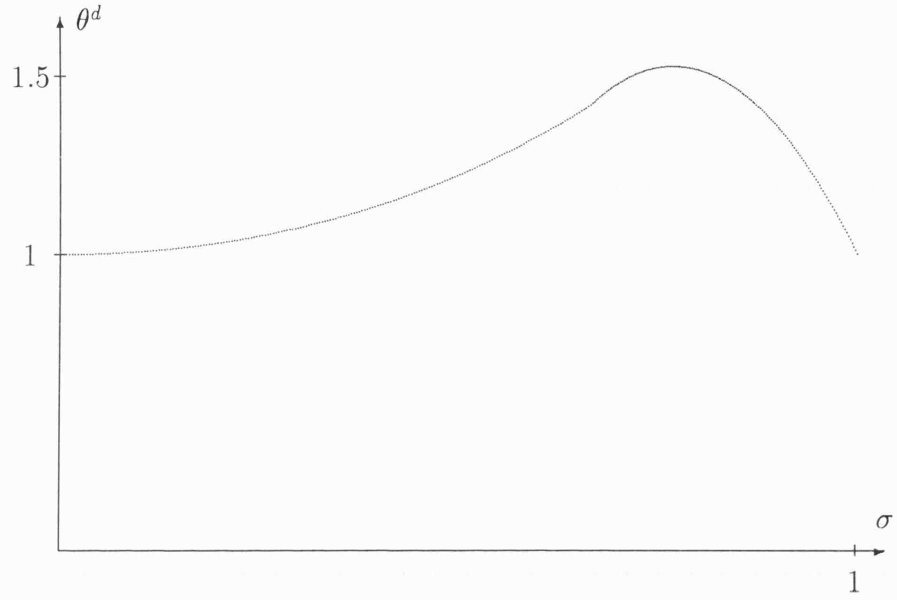

The expressions become considerably simpler in the special case that p^c is the monopoly price $p^m = (a + c)/2$. In this case, the resulting expressions for θ depend only on σ, and we have

$$\theta^d_+ = \frac{(2-\sigma)^2}{4(1-\sigma)} \qquad \text{and} \qquad \theta^d_0 = \frac{(2-\sigma)^2(\sigma + \sigma^2 - 1)}{\sigma^4}. \tag{4.19}$$

The boundary between the two regions is given by $\sigma = \sqrt{3} - 1 \approx 0.732$, and the maximum is attained in the θ^d_0-region at $\sigma \approx 0.775$ at a level of 1.5625. This case is depicted in Figure 4.6. This figure is precisely a mirror image of Figure 2 (right panel) in Deneckere (1983), since the right-hand side $\delta/(1-\delta)$ in (4.7) is the inverse of the discount *rate* corresponding to δ. Thus Figure 4.6 depicts the result familiar from Deneckere (1983) that in this model, there is a nonmonotonic relationship between product differentiation and the stability of collusion: over a broad interval for σ, θ^d is increasing in σ, indicating that stability of collusion is higher if products are more differentiated. For quite similar goods however, where an optimally deviating firm takes over the whole market, this relationship is reversed. As a consequence, for certain values of δ, a particular cartel price may be sustainable for very similar or highly differentiated products, but not for intermediate degrees of product differentiation.

As in the previous model, if δ is sufficiently high (in particular if $\delta > 0.61$, cf. Proposition 4.3 below), then the incentive constraint is satisfied for all values of σ.

The schedule $\theta^d(\sigma)$ has been derived above using the approximated expected demand function q_i^{app} of Lemma 4.3. Since q_1^{app} is an upper bound of the true expected demand of a deviating firm, this means that θ^d, too, is an upper bound. Hence, satisfaction of the incentive constraint (4.7) in terms of the θ^d used here is sufficient but not necessary for (4.7) to be satisfied with respect to true expected demand.

4.5.3 Collusion: marginal deviations

For the probability of a deviation from p^c being detected, we have

Lemma 4.4 *The probability* $\alpha(p_1, p^c) = Prob\{\mathbf{q}(p_1, p^c) \in \Sigma - S(\mathbf{p^c})\}$ *is given by*

$$\alpha(p_1, p^c) = \frac{r(1+\sigma)}{2d(1-\sigma)(r-2\sigma)}(p^c - p_1) + \frac{\sigma(1+r-\sigma)(r-\sigma+\sigma^2)}{4d^2(1-\sigma)^2(r-2\sigma)^2}(p^c - p_1)^2.$$

From Lemma 4.4, the marginal increase in the probability of detection at $p_1 = p^c$ is

$$\frac{\partial\alpha(p^c, p^c)}{\partial p_1} = -\frac{r(1+\sigma)}{2d(1-\sigma)(r-2\sigma)} = -\frac{(1+\sigma)[s(1-\sigma)-2\sigma]}{s(1-\sigma)^2}.$$

Substituting the relevant expressions, the l.h.s. of (4.6) becomes

$$\theta^m = \frac{2ds(1-\sigma)(2-\sigma)^2}{(1+\sigma)[s(1-\sigma)+2\sigma](a+c-2p^c-c\sigma+\sigma p^c)} \tag{4.20}$$

For $\sigma = 1$, $\theta^m = 0$, and for $\sigma = 0$, $\theta^m = 4/(a+c-2p^c)$. The latter result implies that in the special case $p^c = p^m$, θ approaches infinity as $\sigma \to 0$. Moreover, we have

Lemma 4.5 θ^m *is decreasing in* σ *for any value of* s.

This result states that, as in the previous Hotelling model, collusion is *less* sustainable for differentiated goods than for homogeneous goods, as far as marginal

Figure 4.7: Left-hand side of incentive constraint (3.6) for marginal deviations

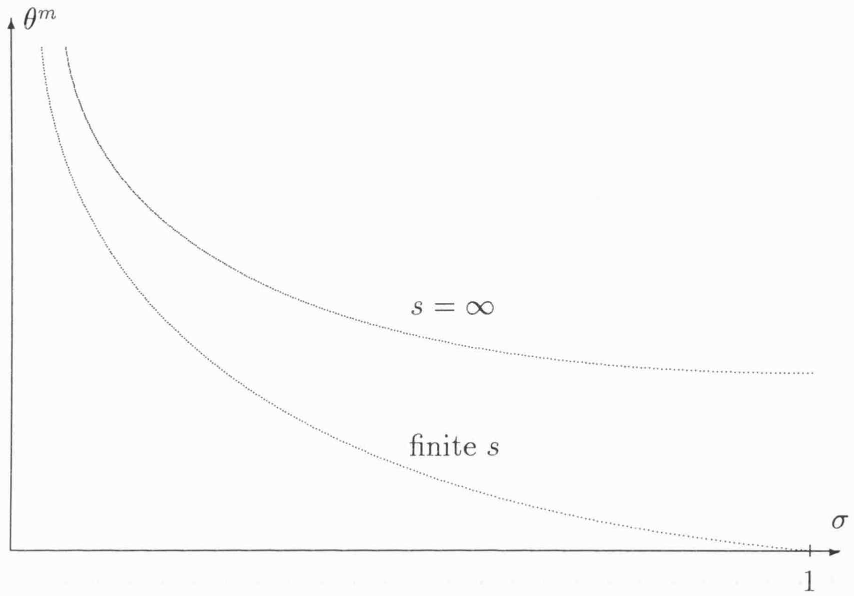

deviations from the cartel price are concerned. The economic intuition for this result has, by and large, already been discussed in Section 4.4.3. A difference, however, arises as to the role of the change of the elasticity of demand as a function of product differentiation: here the two opposite effects of a change in the elasticity described above do not exactly cancel each other. This can be seen by considering the limit case $s \to \infty$, i.e. the case of uncorrelated dmend shocks. It turns out that the schedule $\theta^m(\sigma)$ for this case has the same vertical intercept at $\sigma = 0$ as for any finite value of s, and is decreasing in σ. This latter property implies that the relationship between cartel stability and product differentiation expressed by Lemma 4.5 holds even if the assumption of correlated demand shocks is disposed of. Typical schedules for finite s and for $s \to \infty$ are depicted in Figure 4.7.

Similarly as in the previous example, the correlation effect is reinforced by an increase in the variance of the demand shocks due to differentiation. In order to separate these two effects, one can, instead of considering a fixed noise level d, vary d with σ in such a way that the variance of the demand shocks u_i is always held constant, and then examine the resulting θ^m. It can be shown that the correlation effect is indeed a major force in both this and the previous example. In the Hotelling model, the level of the variance is quite important as well, since it determines the level of the maximum of θ^m for $\sigma = 0$ (cf. the discussion of the role of d in Section 4.4.3). Here, in contrast, it turns out that while changes in the variance of course matter to some extent, the correlation effect is clearly the main force.

Hence in this model, two main forces are present which work in the same direction and imply a lower cartel stability in the case of differentiated goods than for more similar goods: the first is a lower correlation of demand shocks, the second the lower price and cross-price elasticities of demand.

The comparative-statics properties of θ^m with respect to s and d are the same as those of the corresponding expression in the model of Section 4.4: θ^m is decreasing in σ and increasing in both s and d, for quite the same reasons as in Section 4.4.3.

While the results obtained here are broadly similar to the ones obtained for the Hotelling model above, there are three important differences. The first has already been discussed above: the fact that we obtain a negative relationship between product differentiation and cartel stability even without the assumption that the correlation of the demand shocks varies with σ.

Second, the fact that in previous model θ^m does not depend on p^c is indeed a consequence of the Hotelling framework. Here, in contrast, θ^m is increasing in p^c. This means that a decrease in the cartel price relaxes the incentive constraint for marginal deviations. In particular, the incentive constraint (4.6) refers to the sustainability of a particular collusive price, not of collusion as such.

Third, for $p^c = p^m$, θ^m does not have a finite intercept at $\sigma = 0$. This means that for any discount factor $\delta < 1$, there exists a critical level of σ below which the monopoly price cannot be sustained any more.

As in the previous example, before we can bring together the incentive constraints for marginal and for large discrete deviations, we must check whether Condition 4.1 of Section 4.3.4 is satisfied. It is easily verified that this is the

Figure 4.8: Left-hand sides of the incentive constraints (3.6) and (3.7)

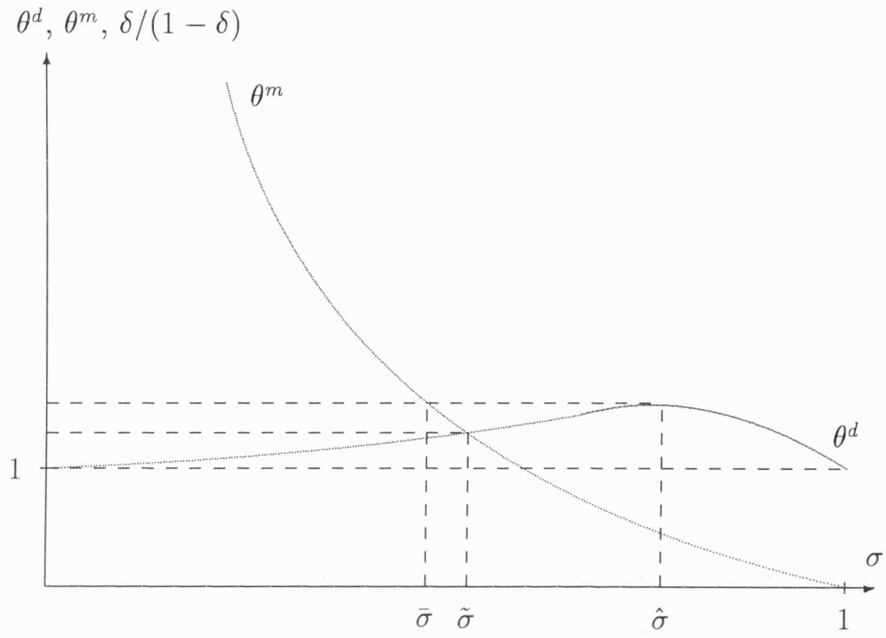

case: π_1 is concave in p_1, and $\alpha(p_1, p^c)$ can be shown to be convex in p_1 as long as $\alpha < 1$. Then, if the general incentive constraint (4.5) is satisfied for marginal deviations, it is satisfied for all prices p_1 such that $\alpha(p_1, p^c) < 1$.

Therefore, we never have to be concerned with small discrete deviations: either (4.6) is satisfied, in which case such deviations are not profitable, or it is not, but then p^c is not sustainable in the first place. Condition 4.1 also solves another potential problem: with the detection probability given by Lemma 4.4, it is not clear that deviation to an "interior" optimal price p^d_+ necessarily implies $\alpha = 1$. If not, such deviations would be more profitable than assumed in the derivation of θ^d. But since Condition 4.1 holds, this precisely is ruled out, at least for the cases where θ^d is relevant, i.e. where (4.6) is satisfied.

4.5.4 Both incentive constraints combined

The left-hand sides of the incentive constraints, θ^d and θ^m, are depicted in Figure 4.8 as functions of σ. As in Section 4.4.4, define $\tilde{\sigma}$ as the σ solving $\theta^d(\sigma) = \theta^m(\sigma)$, and $\delta(\sigma)$ as the solution of the equation $\theta^d(\sigma) = \delta/(1-\delta)$. Moreover let $\hat{\sigma} =$

argmax$\{\theta^d(\sigma)|\sigma \in [0,1]\}$ (cf. Figure 4.8). From Figure 4.8, we then see that there are four different intervals of values of δ:

1. For $\delta < 1/2$, we have $\delta/(1-\delta) < 1$, and p^c is not sustainable for any degree of product differentiation because (4.7) is violated, i.e. it is always profitable for a firm to deviate to the optimal defection price $p^d(p^c)$.

2. For $\delta \in [1/2, \delta(\tilde{\sigma}))$, collusion is sustainable for sufficiently homogeneous goods, but not for more differentiated goods. For larger values of σ, the incentive constraint for discrete deviations is not satisfied, while for lower values of σ, the constraint for marginal deviations is not satisfied.

3. For $\delta \in [\delta(\tilde{\sigma}), \delta(\hat{\sigma})]$, sustainability of p^c depends in a nonmonotonic way on σ: p^c is sustainable for very homogeneous goods and some intermediate levels of differentiation, but not otherwise.

4. Finally, for $\delta > \delta(\hat{\sigma})$, the incentive constraint for discrete deviations is always satisfied. Here, monotonicity of θ^m in σ implies a nonomonotonic relationship between the sustainability of p^c and the degree of differentiation: there exists a critical level of differentiation such that for more homogeneous goods, p^c is sustainable, whereas it is not for more differentiated products. This is more formally stated in the following result:

Proposition 4.3 *For any given collusive price p^c the discount factor $\delta(\hat{\sigma})$ exists and is at most $25/41 \approx 0.61$. For $\delta > \delta(\hat{\sigma})$, there exists $\bar{\sigma}(\delta) \in [0,1)$ such that p^c can (cannot) be sustained for $\sigma \geq$ $(<)$ $\bar{\sigma}(\delta)$. Moreover, if $p^c = p^m$, then $\bar{\sigma}(\delta)$ is strictly positive.*

The economic intuition of this result has already been discussed in Section 4.4. A difference to the previous model is that there, the monopoly price could be sustained for any degree of differentiation, given a sufficiently large discount factor, whereas here, there exists no discount factor such that the monopoly price can be sustained for any σ.

While Proposition 4.3 states only a possibility which depends on the magnitude of the discount factor δ, this case seems to be particularly relevant: according to Proposition 4.3, the critical discount factor $\delta(\hat{\sigma})$ is at most 0.61, i.e. assumes a rather low value (cf. 4.2), whereas in this context of a repeated game with price setting, the relevant discount factor is likely to be rather high (close to one) due to presumably short time intervals.

The comparative-statics properties of $\bar{\sigma}$ immediately follow from the defining equation $\theta^m(\bar{\sigma}) = \delta/(1-\delta)$ and the properties of θ^m discussed above: a decrease in δ implies an increase in $\bar{\sigma}$ for obvious reasons. An increase in s implies an increase in $\bar{\sigma}$ because of the adverse effect of a lower correlation of the demand shocks on the sustainability of p^c. Finally, an increase in d implies an increase in $\bar{\sigma}$ because a noisier environment undermines the sustainability of p^c via a decrease of the detection probability.

4.5.5 Maximal collusive payoff

Until now, we have discussed the stability of collusion in terms of the sustainability of a particular cartel price p^c, e.g. the monopoly price. But in this model (and in general), where the monopoly price is not sustainable, perhaps a lower price might be, since decreasing p^c relaxes both incentive constraints. Consequently, if the notion of "sustainability" of collusion is meaningful only with reference to a particular collusive strategy, "stability" of collusion in general is perhaps more appropriately measured by the maximal collusive payoff (in absolute terms or in comparison with the Bertrand payoff) that can be sustained. We therefore now ask how the maximal sustainable collusive payoff depends on the degree of product differentiation, confining the analysis to the case of Proposition 4.3, i.e. the case of a sufficiently large discount factor.

It is easily verified in (4.20) that θ^m is increasing in p^c. Assuming that δ satisfies the assumption of Proposition 4.3, we then obtain three different intervals for σ:

1. For sufficiently high σ, i.e. sufficiently homogeneous goods, the monopoly price p^m is sustainable.

2. For σ below $\bar{\sigma}(\delta, p^m)$, p^m is not sustainable any more. Here, for given σ, the maximal sustainable price is determined by the incentive constraint (4.6), i.e.

Figure 4.9: Differences $\pi^m - \pi^b$ and $\pi^c_{max} - \pi^b$ as functions of σ

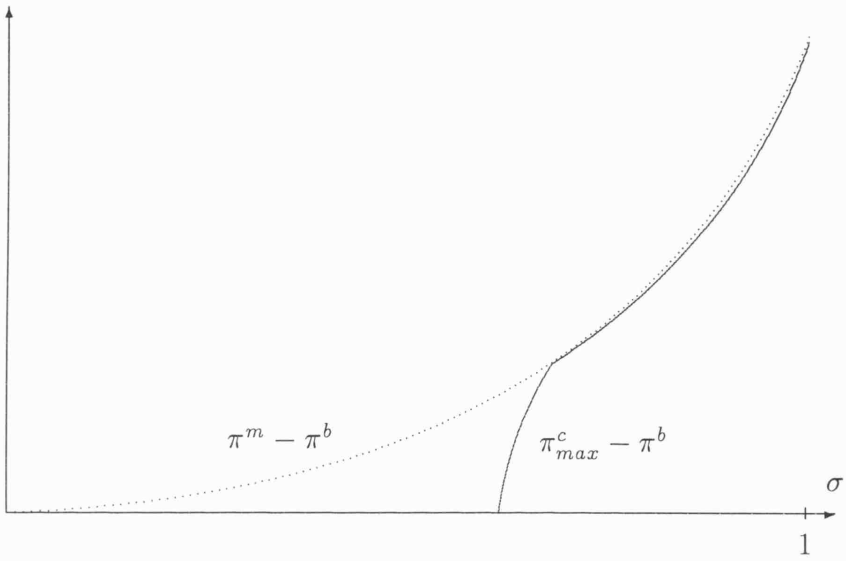

$\theta^m(\sigma, p^c) = \delta/(1-\delta)$, and because of the properties of θ^m, the resulting p^c is an increasing function of σ.

3. But for a critical level of σ, i.e. for a sufficiently large degree of differentiation, the collusive price p^c determined by the above condition reaches the Bertrand price p^b. This means, below a certain level of σ, collusion is not sustainable *at all.* More precisely, due to the monotonicity of θ^m in p^c, sustainability of any price $p^c > p^b$ requires that (4.6) be satisfied for $p^c = p^b$, or

$$\frac{2ds(2-\sigma)^2(1-\sigma)}{\sigma(1+\sigma)(a-c)[s(1-\sigma)+2\sigma]} \leq \frac{\delta}{1-\delta}. \tag{4.21}$$

But with σ in the denominator of (4.21), it follows that there must exist a critical $\sigma > 0$ below which (4.21) is violated.

These results are illustrated in Figure 4.9, which depicts the difference between monopoly and maximal collusive profits, respectively, and the Bertrand payoff. The dotted curve in Figure 4.9 depicts the difference $\pi^m - \pi^b$ and shows a familiar relationship: with p^m constant and a Bertrand price that equals p^m for $\sigma = 0$, i.e. independent goods, and decreases to the level of marginal cost as σ approaches one, this difference is increasing in σ. The solid curve depicts the difference $\pi^c_{max} - \pi^b$. It coincides with the dashed curve as long as p^m is sustainable, but then rapidly drops to the Bertrand price p^b.

Figure 4.9, then, illustrates two points that are related but quite distinct. The first is the well-known fact that with more differentiated goods, collusion is less *profitable* compared to Bertrand competition than with homogeneous goods. The second is the one emphasized in this chapter: with differentiated goods, collusion is not only less profitable, but also less *sustainable.*

These findings complement Kandori's (1992) result that in repeated games with imperfect monitoring, a decrease in the informativeness of a signal implies a decrease in the maximal sustainable payoff. While his model is much more general than the ones considered here, his result only expresses a monotonicity relationship. Here, in contrast, the explicit derivation of optimal strategies is not only technically necessary for the examination of all relevant incentive constraints, but also shows much more clearly how exactly differentiation affects maximal payoffs.

4.6 Comparison of the two examples

In this secion, I will draw together and compare the results obtained in the two examples of Sections 4.4 and 4.5 in order to see which of the results obtained seem to be fairly robust with respect to the properties of the underlying oligopoly models, and which results depend on details of these models. The results common to both models are the following:

1. The incentive constraint for marginal deviations from a cartel price p^c leads to a critical discount factor above which p^c can be sustained, which is a decreasing function of the substitutability parameter σ. In other words, where this constraint is binding, collusion is less stable for more differentiated goods.

2. For lower (actual) discount factors δ, the sustainability of p^c can depend on the degree of product differentiation in a nonmonotonic way. Above a critical level of δ, however (which is 0.5 in Section 4.4 and 0.61 in Section 4.5, i.e. in both cases rather small), the incentive constraint for large deviations is always satisfied.

Here, there exists, at least if the level of uncertainty is sufficiently large, a region of δ where only the incentive constraint for marginal deviations is relevant and leads to the monotonic relationship discussed above: there is a critical $\bar{\sigma}$ above which a cartel price can be sustained and below which it cannot.

3. In both models, for larger degrees of differentation, collusion may not be sustainable at all, i.e. at any price $p^c > p^b$.

Three important differences between the two models should be noted, though:

1. In the Hotelling model, the monotonicity region exists only if there is sufficient demand uncertainty. In contrast, in the representative-consumer model, a monotonicity region exists in any case, and the level of uncertainty determines the critical $\bar{\sigma}$ below which p^c cannot be sustained any more.

2. In the Hotelling model, the assumption that the correlation of the demand shocks depends on σ is essential for a negative relationship between product differentiation and cartel stability. In contrast, in the representative-consumer model, the schedule $\theta^m(\sigma)$ is decreasing in σ even if the two firms' demand shocks are independent.

3. In the Hotelling model, the decrease in cartel stability due to product differentiation is roughly equally driven by a decrease in the correlation of the demand shocks and an increase in their variance. In the representative-consumer model, in contrast, the correlation effect clearly prevails.

4. In the representative-consumer model, "sustainability" of collusion is a meaningful notion only with reference to a particular cartel price p^c: since lowering the collusive price relaxes the incentive constraints, it follows that if the monopoly price cannot be sustained any more, in general a lower price is still sustainable. In the Hotelling model, however, the incentive constraint for marginal deviations does not depend on the particular cartel price under consideration: if it is not satisfied for the monopoly price, then it is not satisfied for any price above p^b. In this model, therefore, it may be meaningful to speak of "sustainability" of collusion *as such.*

In our models, the monotonic relationship emphasized in this chapter, according to which product differentiation adversely affects the stability of collusion, emerges only as a possibility, in addition to cases of nonmonotonicity or even the inverse monotonic relationship derived in other theoretical works. This raises the question how relevant this case is vis-à-vis the other possibilities.

Obviously the level of noise matters: for $d \to 0$, the models converge to deterministic models, in which marginal deviations never occur because they are detected with certainty, an assumption which was criticised in the Introduction as being rather unrealistic. The higher the degree of uncertainty, the more marginal deviations emphasized here become relevant.

Moreover, stepping outside the framework of the two examples, we see that any reason which renders marginal deviations more attractive to a deviant cartel member than large price cuts, implies that collusion is likely to be adversely affected by differentiation. One particularly important factor is limited production capacity: if the current-period profit a deviating firm can gain is limited due to capacity constraints, then large visible price cuts may not be desirable at all, i.e. a firm would rather consider cutting the price only slightly. In this sense, the models presented are biased in favour of large price cuts.

4.7 Concluding remarks

Economists have long believed that product heterogeneity is a factor hindering collusion among oligopolists, because it entails a situation of higher "complexity" than prevails with homogeneous goods. An analytical formulation of this argument, however, has not been available. The purpose of this chapter is to fill this gap in theory, i.e. to suggest a precise formulation of the kind of complexity that heterogeneity brings about.

I have argued that a satisfactory analysis of the stability of collusion should allow that undercutting the cartel price only slightly (in the hope that this is not noticed) be a relevant consideration for the cartel members, which requires a framework with uncertainty and imperfect monitoring. Moreover, uncertainty also provides an essential link between product heterogeneity and collusion: more heterogeneity leads to a decrease in the correlation of the firms' demand shocks. This implies, in a sense, an increase in uncertainty which in turn undermines the stability of collusion.

This effect of product differentiation on the demand system is also likely to be a reason why firms producing heterogeneous products may find it difficult to reach a cartel agreement in the first place.

In contrast, as has been argued in the Introduction, because of the use of very simple deterministic models and the lack of any empirical evidence, the results obtained in the recent theoretical literature do not seem to provide a convincing case for reversing the traditional view.

In his textbook, Tirole (1988, Chap. 5) describes product differentiation and collusion as two possible ways to escape the Bertrand paradox. The results of this chapter, however, suggest that these two solutions are mutually exclusive: where the degree of differentiation is a choice parameter, firms can either raise margins by differentiating their products, or seek to abolish price competition, in which case products must be *standardized.*

Because of the technical difficulties the explicit derivation of optimal collusive strategies entails, I have illustrated the above ideas using very specific oligopoly models, restricting the collusive strategies to a plausible yet ad-hoc class, and making specific assumptions on the distribution of the demand shocks. It remains to be seen to what extent the results obtained here hold more generally. At least as far as the choice of the particular oligopoly model is concerned, the similarity of the main predictions obtained in two quite different models suggests that the economic forces that link the stability of collusion to the degree of product differentiation via uncertainty are reasonably robust.

4.8 Appendix: Proofs

Proof of Proposition 4.1: Consider a collusive strategy characterized by the triple $s^c = (p^c, B, T)$. Assuming that firm 2 adheres to the cartel price p^c, then firm 1's period profit is $\pi(p_1, p^c)$, and define $\alpha(\mathbf{p}) = \text{Prob}\{\mathbf{q}(\mathbf{p}) \in \Sigma - S(\mathbf{p^c})\}$ and $\beta(\mathbf{p}) = \text{Prob}\{\mathbf{q}(\mathbf{p}) \in B\}$. Then firm 1's expected discounted payoff is characterized by the Bellman equation

$$\begin{aligned} v_1(p_1, p^c) &= \pi_1(p_1, p^c) + [1 - \alpha(p_1, p^c) - \beta(p_1, p^c)]\delta v_1(p_1, p^c) + \alpha(p_1, p^c)\frac{\delta}{1-\delta}\pi^b \\ &+ \beta(p_1, p^c)\delta\left(\sum_{\tau=0}^{T-2}\delta^\tau \pi^b + \delta^{T-1}v_1(p_1, p^c)\right), \end{aligned} \tag{4.22}$$

where π^b is the period profit in the static Bertrand equilibrium. Here, the third term on the r.h.s. specifies the payoff obtained if a deviation can be inferred with certainty, in which case the firms play the Bertrand equilibrium forever. The fourth term specifies the payoff obtained if a quantity vector in B occurs, in which case a price war of duration $T-1$ follows (cf. Green/Porter 1984). Equation (4.22) can be solved explicitly for v_1:

$$v_1 = \frac{\pi^b}{1-\delta} + \frac{\pi_1 - \pi^b}{1 - \delta + \delta\alpha + (\delta - \delta^T)\beta}, \tag{4.23}$$

where the arguments of v_1, π_1, α, and β have been omitted.

Let η_0 be defined as in the proposition (it exists since by assumption B is compact). Suppose now that B in the collusive strategy does not satisfy the condition $\mathbf{q} \in B$ for all $\mathbf{q}$ for which $\eta(\mathbf{q}) < \eta_0$ (here and in the following, zero-measure sets are ignored). Then there exists some $\tilde{\eta} \leq \eta_0$ such that we can select two arbitrary subsets with positive measure: $C_1 \subset B$ such that $\eta(\mathbf{q}) \geq \tilde{\eta}\ \ \forall \mathbf{q} \in C_1$ and $C_2 \subset S(\mathbf{p^c}) - B$ such that $\eta(\mathbf{q}) < \tilde{\eta}\ \ \forall \mathbf{q} \in C_2$. Given these sets, let $C_0 = B - C_1$. We show that in this case s^c cannot be optimal because it is dominated either by a strategy with the trigger set C_0 or one with the set $B \cup C_2$. Thus, the proof works as follows: If the condition of the proposition is not met, then B fails to be optimal for either of two reasons: either it is not desirable to have C_1 contained in B in the first place, or if it is, then it must be even better to include the set C_2 as well, given the assumptions on C_1 and C_2. More specifically, we consider two alternative strategies $s_0^c = (p^c, C_0, T_0)$ and $s_2^c = (p^c, C_0 \cup C_1 \cup C_2, T_2)$ which are sustainable and yield the same payoff as s^c. For the latter codition to hold, it follows from (4.23) that T_0 and T_2 must be chosen such that

$$(\delta - \delta^{T_0})\beta_0 = (\delta - \delta^{T_2})\beta_2 = (\delta - \delta^T)\beta, \tag{4.24}$$

where β_0 and β_2 are the price war probabilities corresponding to the trigger sets of s_0^c and s_2^c, respectively.

For s^c to be sustainable, marginal deviations from p^c must not be profitable. Differentiation of (4.23) with respect to p_1 leads to the condition

$$\frac{\partial \pi_1}{\partial p_1}\left[1 - \delta + (\delta - \delta^T)\beta\right] - \left[\delta \frac{\partial \alpha}{\partial p_1} + (\delta - \delta^T)\frac{\partial \beta}{\partial p_1}\right] \geq 0, \tag{4.25}$$

which by assumption is binding for s^c. Similar conditions, not necessarily binding, must hold for s_0^c and s_2^c to be sustainable. Given the equality in payoffs of s^c, s_0^c and s_2^c, it follows that s^c is not optimal if the condition (4.25) is not binding for one of the two alternative strategies, because then the sustainable payoff could be increased by choosing a different trigger set. By construction of the strategies, the only relevant terms in (4.25) to compare are the $(\delta - \delta^T)(\partial\beta/\partial p_1)$-terms, which are always negative. Hence, optimality of B requires

$$(\delta - \delta^T)\frac{\partial \beta}{\partial p_1} \leq \min\left\{(\delta - \delta^{T_0})\frac{\partial \beta_0}{\partial p_1}, (\delta - \delta^{T_2})\frac{\partial \beta_2}{\partial p_1}\right\},$$

or, using (4.24),

$$\frac{\partial \beta}{\partial p_1} \leq \min\left\{\frac{\partial \beta_0}{\partial p_1}\frac{\beta}{\beta_0}, \ \frac{\partial \beta_2}{\partial p_1}\frac{\beta}{\beta_2}\right\}. \tag{4.26}$$

For a more compact notation, write x_i for $\mathrm{Prob}\{\mathbf{q} \in C_i\}$, and y_i for $\frac{\partial}{\partial p_i}\mathrm{Prob}\{\mathbf{q} \in C_i\}$, for $i \in \{0,1,2\}$. Now it is straightforward to show that by construction of the sets C_1 and C_2, we have $y_2/x_2 < y_1/x_1$, and eqn. (4.26) implies the inequalities

$$y_0 + y_1 \leq y_0 \frac{x_0 + x_1}{x_0} \quad \text{and} \quad y_0 + y_1 \leq (y_0 + y_1 + y_2)\frac{x_0 + x_1}{x_0 + x_1 + x_2}. \tag{4.27}$$

Then the first of these inequalities together with the preceding inequality implies $y_2/x_2 < y_1/x_1 \leq y_0/x_0$, which leads to a contradiction with the second inequality in (4.27). This proves the first part of the proposition.

As for the second part, consider a strategy s^c that satisfies condition 1. of the Proposition, but assume that both sets

$$C_1 = \{\mathbf{q}|\eta(\mathbf{q}) = \eta_0\} \cap B \quad \text{and} \quad C_2 = \{\mathbf{q}|\eta(\mathbf{q}) = \eta_0\} - B$$

have nonzero measure. Now we run again through the proof of part 1: By construction, $y_2/x_2 = y_1/x_1$, and there are two possibilities: If $y_1/x_1 < y_0/x_0$, we again arrive at the same contradiction. This means that if it is strictly better to include C_1 in B rather than to exclude it, it is still better to include C_2 as well, which is the statement of part 2a). Otherwise, $y_1/x_1 = y_0/x_0$, and we have $y_2/x_2 = y_1/x_1 = y_0/x_0$. Here, the three strategies are equivalent, since the incentive constraints and the resulting payoffs are the same. This, however, means that the subset $\{\mathbf{q}|\eta(\mathbf{q}) = \eta_0\}$ cannot be relevant at all, since by continuity a similar strategy with a trigger set $B - \{\mathbf{q}|\eta(\mathbf{q}) = \eta_0\}$ would lead to the same payoff. ■

Proof of Proposition 4.2: Given the uniform distribution of $\mathbf{v}$ on $[\mu - d, \mu + d]^2$ (cf. Section 3.1) and the fact for any price vector, there is one-to-one mapping from $\mathbf{v}$ into $\mathbf{q}$, it follows that all demand vectors $\mathbf{q} \in S(\mathbf{p})$, i.e. those that can be observed for a given price vector $\mathbf{p}$, have the same density and that also the relative likelihood change $(\partial f/\partial p_i)/f$ is the same for all vectors in the interior of $S(\mathbf{p})$.

Then from Proposition 4.1, part 2., it follows that if it were desirable to include any vectors of $S(p^c, p^c)$ in the trigger set B, then all of them should be included, which clearly cannot be optimal. Therefore, this cannot be the case. The only remaining possibilities are that a trigger set $B \subset S(\mathbf{p^c})$ is either not payoff-relevant at all (which is quite unlikely, but not excluded by Proposition 4.1), or that it is strictly undesirable to have a nonempty trigger set B. In either case, setting $B = \emptyset$ is optimal here. ■

Proof of Lemma 4.1:

From the derivatives of $\mathbf{q}$ and $\mathbf{u}$ with respect to $\mathbf{v}$ it follows that the supports of $\mathbf{q}$ and $\mathbf{u}$ in their respective spaces are rhombuses enclosed by the lines corresponding to $v_1 = 1-d$, $v_1 = 1+d$, $v_2 = 1-d$ and $v_2 = 1+d$. For any fixed v_2^0, the points $\mathbf{u}$ on the line $v_2 = v_2^0$ are characterized by $(\sigma/r)u_1 - ((r-\sigma)/r)u_2 = -((r-2\sigma)/r)v_2^0$. With (4.9), it then follows that if $p_1 = p^c$, then

$$(\sigma/r)q_1 - ((r-\sigma)/r)q_2 = -\frac{r-2\sigma}{2r}v_2 \in [-\frac{r-2\sigma}{2r}(1+d), -\frac{r-2\sigma}{2r}(1-d)].$$

Given that if firm 1 is the potential deviant firm, the latter expression is always at least $-(r-\sigma)(1+d)/(2r)$, but it can exceed the upper bound if $pl < p^c$, i.i $k < 0$. Therefore, a price war is triggered if $\sigma q_1/r - (r-\sigma)q_2/r > (r-2\sigma)(1-d)/(2r)$, or

$$k(r-\sigma)v_2 + (r-2\sigma)tv_2 < (r-2\sigma)t(1-d) - k\sigma v_1, \tag{4.28}$$

which describes a positively sloped line in (v_1, v_2)-space. By a similar argument, the points $\mathbf{u}$ on the line $v_1 = v_1^0$ are for $k = 0$ characterized by

$$\frac{r-\sigma}{r}q_1 - \frac{\sigma}{r}q_2 = \frac{r-2\sigma}{2r}v_1 \leq \frac{r-2\sigma}{2r}(1+d).$$

Thus we have an upper bound on $(r-\sigma)\tilde{q}_1/r - \sigma\tilde{q}_2/r$, and a price war is triggered if

$$[(r-2\sigma)t - k\sigma]v_1 - k(r-\sigma)v_2 > (r-2\sigma)(1+d)t, \tag{4.29}$$

which describes a negatively sloped line in (v_1, v_2)-space. Given the lines (4.28) and (4.29), the price war probability can be calculated in the following way. Let $v_1^i, v_2^i, \underline{v}_1, \underline{v}_2$ be defined as illustrated in Figure 4.10. Then α can be calculated as

$$\begin{aligned}\alpha &= 1 - \frac{1}{4d^2}\Big[(1+d-v_2^i)(\underline{v}_1 - 1 + d) + \frac{1}{2}(v_1^i - 1 + d)(v_2^i - \underline{v}_2) \\ &\quad + \frac{1}{2}(1+d-v_2^i)(v_1^i - \underline{v}_1)\Big],\end{aligned}$$

which leads to the expression stated in the Lemma. ■

Proof of Lemma 4.2 Using the relevant expressions for $(\pi_1(k,p^c) - \pi^c)/(\pi^c - \pi^b)$ and $\alpha(k,p^c)$, the condition $\partial^2\pi(k,p^c)/\partial k^2 \leq \partial^2\alpha(k,p^c)/\partial k^2$ at $k = 0$ is equivalent to

$$\frac{2}{p^c - t} \geq \frac{8d^2\sigma^2 + r(1-d)(r - 3rd + 8d\sigma)}{2d^2(r-2\sigma)^2 t}. \tag{4.30}$$

Figure 4.10: Calculation of α in the Hotelling model

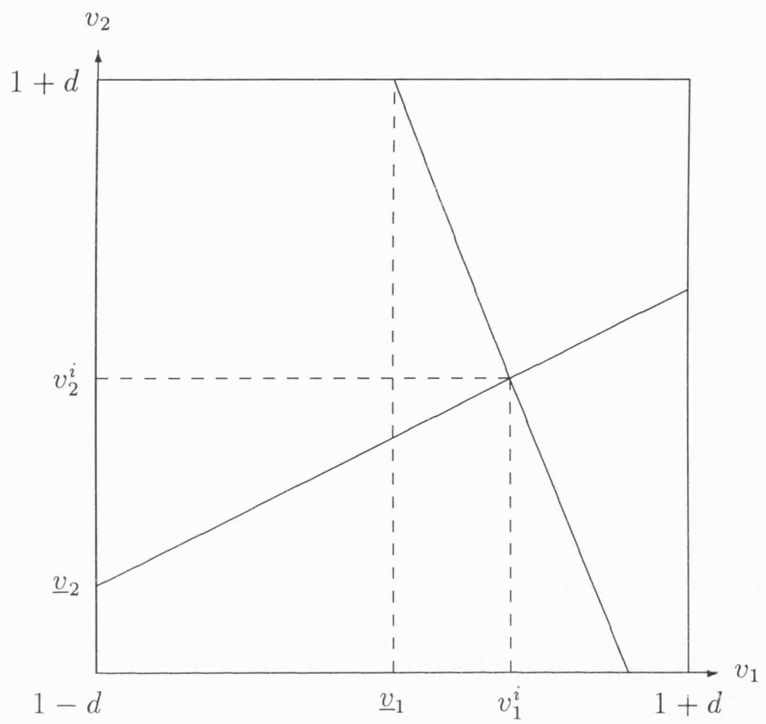

By assumption, (4.6) is binding, and the value of σ as a function of β is given by

$$\sigma(\beta) = \frac{(2+\beta)d - 2\beta}{4\beta - 2\beta s + 2ds + \beta ds}.$$

Substitution of $\sigma(\beta)$ then leads to (4.12). The upper bound in (4.12) can be shown to be a negative function of d. Therefore, if (4.12) is satisfied for $d = 1$, it is satisfied for all d, and substituting $d = 1$ into (4.12) leads to (4.13). Finally, the upper bound in (4.13) can be shown to be a positive function of β. Thus, if $\delta \geq 1/2$, then $\beta \geq 1$, and substituting $\beta = 1$ into (4.13) leads to the simpler upper bound stated in the Lemma. ■

Proof of Lemma 4.3: From (4.15) and (4.16), we have $q_1 = q_1^n$ for $u_2 \geq \bar{u}_2$ and $q_1 = q_1^m$ for $u_2 \leq \bar{u}_2$, where

$$\bar{u_2} = \frac{1}{1-\sigma}\left[p^c - \sigma p^1 - (1-\sigma)a\right]$$

is the value of u_2 where $q_2^n = 0$. We show that the expected value of q_1 can be at most $q_1^n(\mathbf{p},\mathbf{0})$. Ignoring the constant factor m, we have

$$\begin{aligned}
Eq_1(\mathbf{p},\mathbf{u}) &= \text{Prob}\{u_2 \geq \bar{u_2}\}\left[a + E(u_1|u_2 \geq \bar{u_2}) + \frac{\sigma p_2 - p_1}{1-\sigma}\right] \\
&\quad + \text{Prob}\{u_2 \leq \bar{u_2}\}\left[(1+\sigma)(a-p_1) + E(u_1 + \sigma u_2 | u_2 \leq \bar{u_2})\right] \\
&\leq \text{Prob}\{u_2 \geq \bar{u_2}\}\left[a + E(u_1|u_2 \geq \bar{u_2}) + \frac{\sigma p_2 - p_1}{1-\sigma}\right] \\
&\quad + \text{Prob}\{u_2 \leq \bar{u_2}\}\left[a + E(u_1|u_2 \leq \bar{u_2}) + \frac{\sigma p_2 - p_1}{1-\sigma}\right] \\
&= a + \frac{\sigma p_2 - p_1}{1-\sigma} = q_1^n(\mathbf{p},\mathbf{0}).
\end{aligned}$$

Here, the inequality follows from the fact that $u_2 \leq \bar{u}_2$ is equivalent to

$$(1+\sigma)(a-p_1) + u_1 + \sigma u_2 \leq a + u_1 + \frac{\sigma p_2 - p_1}{1-\sigma},$$

and the equality results from evaluating the second expression, where

$$\text{Prob}\{u_2 \geq \bar{u}_2\}E(u_1|u_2 \geq \bar{u}_2) + \text{Prob}\{u_2 \leq \bar{u}_2\}E(u_1|u_2 \leq \bar{u}_2) = E(u_1) = 0.$$

An analogous argument shows that $Eq_1(\mathbf{p},\mathbf{u}) \leq q_1^m(\mathbf{p},\mathbf{0})$, which completes the proof. ∎

Proof of Lemma 4.4:

Define $\tilde{q}_i := q_i/m - a + p^c$. Then from (4.15) we have

$$\tilde{q}_1 = u_1 + \frac{p^c - p_1}{1-\sigma} \qquad \text{and} \qquad \tilde{q}_2 = u_2 - \frac{\sigma(p^c - p_1)}{1-\sigma}. \tag{4.31}$$

The proof follows the same arguments as the proof of Lemma 4.1. With (4.31), it follows that if $p_1 = p^c$, then $(\sigma/r)\tilde{q}_1 - ((r-\sigma)/r)\tilde{q}_2$ must lie in the interval $\pm((r-\sigma)/r)d$. Given that if firm 1 is the potential deviant firm, a price war is then triggered if $(\sigma\tilde{q}_1)/r - (r-\sigma)\tilde{q}_2/r > (r-\sigma)d/rd$ or

$$v_2 < -d + \frac{\sigma(1+r-\sigma)}{(1-\sigma)(r-2\sigma)}(p^c - p_1). \tag{4.32}$$

Similarly, a price war is triggered if $(r-\sigma)\tilde{q}_1/r - \sigma\tilde{q}_2/r > ((r-2\sigma)/r)d$ or

$$v_2 < -d + \frac{\sigma(1+r-\sigma)}{(1-\sigma)(r-2\sigma)}(p^c - p_1). \tag{4.33}$$

With v_1 and v_2 uniformly distributed and independent, the probability of the event $\{(4.32)$ or $(4.33)\}$ is straightforward to calculate and leads to the expression stated in the Lemma. ∎

Proof of Lemma 4.5: With θ^m as stated in (4.20), $\partial\theta^m(\sigma)/\partial\sigma$ can be expressed as $-x[s(1-\sigma)^2y+z]$, where x is a positive factor and z is not a function of s. It can then be shown that y is positive since $p^c \in [p^c, p^m]$. For z, we have $\partial z/\partial p^c < 0$, and since $z > 0$ for $p^c = p^m$ it follows that z must be positive for all p^c. ∎

Proof of Proposition 4.3: First, since θ^d_+ is increasing in σ, it follows that the maximum of $\theta^d(\sigma)$ must be in the θ^d_0-region. Here, it can be shown that for any σ, θ^d_0 for $p^c = p^m$ is always greater than or equal to θ^d for any lower p^c (this entails looking at the difference of these expressions rather than the derivative with respect to p^c). It follows that $\delta(\hat{\sigma})$ for any p^c never exceeds $\delta(\hat{\sigma})$ for $p^c = p^m$. Using (4.19), the maximum of θ^d is easily calculated as 25/16, from which the upper bound 25/41 for $\delta(\hat{\sigma})$ follows. The second part of the Proposition is an immediate consequence of Lemma 4.5. ∎

References

Abreu, Dilip; David Pearce and Ennio Stacchetti 1986: Optimal cartel equilibria with imperfect monitoring, *Journal of Economic Theory* 39: 251-269

Abreu, Dilip; David Pearce and Ennio Stacchetti 1990: Towards a theory of discounted repeated games with imperfect monitoring, *Econometrica* 58: 1041-1064

Archibald, G.C; B.C. Eaton and R.G. Lipsey 1986: Address models of value theory, in: J.E. Stiglitz and G.F. Mathewson (eds.), *New developments in the analysis of market structure* (MIT), 3-47

Basar, Tamer and Yu-Chi Ho 1974: Informational properties of the Nash solutions of two stochastic nonzero-sum games, *Journal of Economic Theory* 7: 370-387

Chang, Myong-Hun 1991: The effects of product differentiation on collusive pricing; *International Journal of Industrial Organization* 9: 453-469

Clarke, Richard 1983: Collusion and the incentives for information sharing, *Bell Journal of Economics* 14: 383-394

Clarke, Roger 1985: Industrial Economics (Basil Blackwell)

Claycombe, Richard J. 1991, Spatial retail markets, *International Journal of Industrial Organization*, 9: 303-313

Claycombe, Richard J. and Tamara E. Mahan 1993, Spatial aspects of retail market structure - Beef pricing revisited, *International Journal of Industrial Organization*, 11: 283-291

d'Aspremont, Claude, Jean Jaskold Gabszewicz and Jacques-François Thisse 1979: On Hotelling's Stability in Competition, *Econometrica*, 47: 1045-1050

de Palma, André; Victor Ginsburgh, Y. Y. Papageorgiou and Jacques-François Thisse 1985: The principle of minimal differentiation holds under sufficient heterogeneity, *Econometrica* 53: 767-781

de Palma, André; Robin Lindsey; Balder von Hohenbalken; Douglas S. West 1994: Spatial price and variety competition in an urban retail market - a nested logit analysis, *International Journal of Industrial Organization* 12: 331-357

Deneckere, R. 1983: Duopoly supergames with product differentiation, *Economics Letters* 11: 37-42

Eaton, B. Curtis and Richard G. Lipsey 1975: The principle of minimal differentiation reconsidered: some new developments in the theory of spatial competition, *Review of Economic Studies* 42: 27-49

Eaton, B. Curtis and Richard G. Lipsey 1982, An economic theory of central places, *Economic Journal* 92: 56-72

Fraas, Arthur G.; and Douglas Greer 1977: Market structure and price collusion: an empirical analysis, *Journal of Industrial Economics* 26: 21-44

Fried, Dov 1984: Incentives for information production and disclosure in a duopolistic environment, *Quarterly Journal of Economics* 99: 367-381

Friedman, J.W.; and J.F. Thisse 1993: Partial collusion yields minimal product differentiation, *RAND Journal of Economics*

Fudenberg, Drew; and David Levine 1991: An approximate folk theorem with imperfect private information, *Journal of Economic Theory* 54:26-47

Gabszewicz, Jean Jaskold and Jaques-François Thisse 1986: Spatial competition and the location of firms, in J. J. Gabszewicz, J.-F. Thisse, M. Fujita, U. Schweizer (eds.), *Fundamentals of pure and applied economics. Vol. 5: Location theory*, Chur, Switzerland: Harwood, 1-71

Gal-Or, Esther 1985: Information sharing in oligopoly, *Econometrica* 53: 329-343

Gal-Or, Esther 1986: Information transmission - Cournot and Bertrand equilibria, *Review of Economic Studies* 53: 85-92

Green, Edward J.; and Robert H. Porter 1984: Noncooperative collusion under imperfect price information, *Econometrica* 52: 87-100

Häckner, Jonas 1993: Product differentiation and the sustainability of collusion, in: Häckner, Price and quality - essays on product differentiation, Stockholm (IUI Dissertation Series)

Häckner, Jonas 1994: Collusive pricing in markets for vertically differentiated products, *International Journal of Industrial Organization* 12: 155-177

Hay, George A.; and Daniel Kelley: An empirical survey of price-fixing conspiracies, *Journal of Law and Economics* 17: 13-38

Hotelling, Harold 1929: Stability in competition, *Economica* 39: 41-57

Hviid, Morten 1989: Risk-averse duopolists and voluntary information transmission, *Journal of Industrial Economics* 38: 49-64

Hwang, Hae-Shin 1993: Optimal information acquisition for heterogeneous duopoly firms, *Journal of Economic Theory* 59: 385-402

Jehiel, Philippe 1992: Product differentiation and price collusion; *International Journal of Industrial Organization*

Jin, Jim 1992: Information sharing in oligopoly: a general model, mimeo, Wissenschaftszentrum Berlin

Kandori, Michihiro 1992: The use of information in repeated games with imperfect monitoring, *Review of Economic Studies* 59: 581-593

Kirby, Alison J. 1988: Trade associations as information exchange mechanisms, *RAND Journal of Economics* 19: 138-146

Klemperer, Paul 1992: Equilibrium product lines: competing head-to-head may be less competitive, *American Economic Review* 82: 740-755

Kreps, David M.; and José Scheinkman 1983: Quantity precommitment and Bertrand competition yield Cournot outcomes, *Bell Journal of Economics* 14: 326-337

Lehrer, Ehud 1992: On the equilibrium payoffs set of two player repeated games with imperfect monitoring, *International Journal of Game Theory* 20: 211-226

Li, Lode; Richard D. McKelvey and Talbot Page 1987: Optimal research for Cournot oligopolists, *Journal of Economic Theory* 42: 140-166

Li, Lode 1985: Cournot oligopoly with information sharing, *RAND Journal of Economics* 16: 521-536

Madansky, Albert 1986: Foundations of econometrics, Amsterdam and Oxford

Martin, Stephen 1993: Advanced industrial economics (Blackwell)

Neven, D.; and J.-F. Thisse 1990: On quality and variety competition; in: J. J. Gabszewicz, J.-F. Richard, and L. Wolsey (eds.), *Economic decision making: games, econometrics, and optimization. Contributions in honour of Jacques H. Drèze*, Amsterdam: North-Holland, pp. 175-199

Novshek, William and Hugo Sonnenschein 1982: Fulfilled expectations Cournot duopoly with information acquisition and release, *Bell Journal of Economics* 13: 214-218

Porter, Robert 1983: Optimal cartel trigger-price strategies, *Journal of Economic Theory* 29: 313-338

Posner, Richard A. 1976: Antitrust law - an economic perspective, Chicago and London (Univ. of Chicago Press)

Radner, Roy 1962: Team decision problems, *Annals of Mathematical Statistics* 33: 857-81

Raith, Michael 1993: A general model of information sharing in oligopoly, STICERD Discussion Paper TE/93/260

Ross, Thomas W. 1992: Cartel stability and product differentiation, *International Journal of Industrial Organization* 10:1-13

Sakai, Yasuhiro 1985: The value of information in a simple duoppoly model, *Journal of Economic Theory* 36: 36-54

Sakai, Yasuhiro 1986: Cournot and Bertrand equilibria under imperfect information, *Journal of Economics (Zeitschrift für Nationalökonomie)* 46: 213-32

Sakai, Yasuhiro and Takehiko Yamato 1989: Oligopoly, information and welfare, *Journal of Economics (Zeitschrift für Nationalökonomie)* 49: 3-24

Salop, Steven C. 1979: Monopolistic competition with outside goods, *Bell Journal of Economics*, 10:141-156

Scherer, F.M.; and David Ross 1990: Industrial market structure and economic performance (Houghton Mifflin)

Shapiro, Carl 1986: Exchange of cost information in oligopoly, *Review of Economic Studies* 53: 433-446

Shapiro, Carl 1989: Theories of oligopoly behavior, in: Richard Schmalensee and Robert Willig (eds.), *Handbook of Industrial Organization*, Amsterdam (North-Holland), 329-414

Stahl, Konrad 1987, Theories of urban business location, in: Mills (ed.), *Handbook of regional and urban economics, Vol. II*, Elsevier

Stigler, George 1964: A theory of oligopoly, *Journal of Political Economy* 72:44-61

Sultan, Ralph G.M. 1974: Pricing in the electrical oligopoly, Vol. I, Cambridge/Mass. (Harvard)

Tedeschi, Piero 1994: Cartels with homogeneous and heterogeneous and imperfect monitoring, *International Economic Review* 35: 635-656

Theil, Henri 1971: Principles of econometrics, Santa Barbara et al.

Thill, Jean-Claude and Isabelle Thomas 1987, Toward conceptualizing trip-chaining behavior: a review, *Geographical Analysis*, 1-17

Tirole, Jean 1988: The theory of industrial organization, Cambridge/Mass. and London (MIT)

Ulen, Thomas 1983: Railroad cartels before 1887: the effectiveness of private enforcement of collusion, *Research in Economic History* 8: 125-144

Verboven, Frank 1995: Localized competition, multimarket operation and collusive behavior, mimeo, CentER Tilburg

Vives, Xavier 1984: Duopoly information equilibrium: Cournot and Bertrand, *Journal of Economic Theory* 34: 71-94

Vives, Xavier 1988: Aggregation of information in large Cournot markets, *Econometrica* 56: 851-876

Vives, Xavier 1990: Trade association disclosure rules, incentives to share information, and welfare, *RAND Journal of Economics* 21: 409-430

Wernerfelt, Birger 1989: Tacit collusion in differentiated Cournot games, *Economics Letters* 29: 303-306

Zhang, Z. John 1995: Price-matching policy and the principle of minimum differentiation, *Journal of Industrial Economics* 43: 287-299

Zvi, Amir 1993: Information sharing in oligopoly: the truthtelling problem, *RAND Journal of Economics* 24: 455-465