Optimal product-line decisions in dynamic markets

Antonio Buttà

Submitted for the degree of PhD

London School of Economics

September 2007

I certify that the thesis I have presented for examination for the PhD degree of the London School of Economics and Political Science is solely my own work.

I warrant that this authorization does not, to the best of my belief, infringe the rights of any third party.

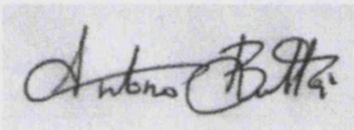

Abstract

In differentiated markets, product-line decisions are a key dimension of the competitive process. In this thesis, I study firms' product-line choices in dynamic vertically-differentiated markets, capturing the process of "creative destruction" whereby technological progress allows firms to introduce products of better quality in the market that affect the profitability and ultimately the viability of relatively obsolete products of lower quality.

The starting point of my analysis is the development of a static empirical structural model of endogenous quality choice, which I use to study the determinants of optimal product lines in the market for personal computers. I empirically estimate the structure of this model and I use it to study the relationship between consumers' willingness to pay for quality and the equilibrium range of qualities sold in the market. I then take a closer look at firms' product-line choices by developing an empirical structural dynamic model of optimal product-line decisions in the presence of sunk adjustment costs. This model allows me to capture the relative rigidity of product lines that I observe in the market for personal computers, to estimate firms' product-line adjustment costs and to study the impact that these costs have on firms' profitability.

I also illustrate two theoretical studies that provide additional insights on firms' optimal product-line decisions in innovative markets. The first study

investigates the intra-firm cannibalisation considerations that underlie the process of product replacement. I then analyse the product-line decisions taken by durable-good manufacturers and, in particular, I present a model to explain why durable good manufacturers in dynamic industries often play an active role as intermediaries in the exchange of second-hand, relatively obsolete, goods.

Contents

List of Tables

List of Figures

Acknowledgements

I would like to acknowledge the advice and support of John Sutton, Peter Davis, Mark Schankerman, Philipp Schmidt-Dengler, Thomas Buettner, Alberto Salvo-Farre and Katharina Sailer. I am particularly grateful to Christos Genakos, for the long discussions we had, for the dataset that he generously shared with me and, above all, for his contagious passion for economics.

This dissertation would never have been written if it had not been for the encouragement, support and assistance I received from my family, friends and colleagues, many of whom I will not have the opportunity to thank individually.

First of all, I would like to thank Cecilia for giving me the time and the space to do what I had to do.

I am also utterly grateful to many dear friends who have been my family in London and in particular to Anì, Barbara, donnaEmla, Giancarlo, Marko and ziaRen. Ciccio has always been there, despite the distance, as only a true friend could have been.

Studying in London would have not been possible without the help of my family. They have always respected and supported my choices, even the difficult ones.

I would also like to thank Europe Economics and Ofcom for the opportunity to escape from my research and to work as an economist in the real world.

I have learnt a lot from my colleagues, some of whom have also become dear friends.

Introduction

In differentiated markets, product-line choices are a key dimension of the competitive process. In this thesis, I present four empirical and theoretical studies of optimal product-line decisions.

In Chapter 1, I develop and estimate a static structural model of endogenous quality choice, which I use to study the determinants of optimal product lines in the market for personal computers. This is a prototypical example of a dynamic vertically-differentiated market in which the range of qualities sold rapidly evolves over time as ongoing innovation allows firms to introduce products of better quality in their product lines and obsolete products of lower quality are withdrawn from the market in a relatively ordered sequence.

In this model, I explain the equilibrium range of product-qualities sold in the market as the result of optimal product-line decisions adopted by a set of monopolistically competitive firms given exogenous, though evolving, technological conditions. I recover empirically the structure of demand and supply of this entry model and I simulate the impacts that a change in consumers' willingness to pay for quality and a change in the cost of providing product variety have on the range of product-qualities sold in the market and on firms' profitability.

In Chapter 2, I also follow a structural empirical approach to the study of quality choice. In particular, I extend the static model presented in Chapter

1 to a dynamic specification, accounting for the sunk costs that firms incur when modifying their product lines. The existence of (fixed) sunk costs implies that product-line decisions have a dynamic nature and that firms adjust their product lines discontinuously, only when the benefits of doing so exceed the costs. This feature of the model matches well the relative rigidity of product-lines that I observe in the market for personal computers.

I estimate the structure of this dynamic model by using the approach developed by Bajari, Benkard and Levin (2006) and, in particular, I derive estimates of the costs that firms incur when modifying their product lines. To the best of my knowledge, this is the first study that empirically estimates product-line adjustment costs.

I use the calibrated model to simulate the impact of a reduction of product-line adjustment costs on firms' profitability.

From the perspective of an individual firm, the existence of fixed sunk adjustment costs has both a direct and an indirect negative impact on a firm's profitability. The direct impact is the outlay that the firm incurs every time it modifies the product line. The indirect effect on profitability stems instead from the impact that adjustment costs have on the product line optimally sold by the firm. When changes to product lines entail a significant fixed sunk cost, the firm will not optimally modify the products it sells every period to adapt to changing technological and market conditions, but only when the (expected) benefits of doing so exceed the costs. This product-line rigidity implies that, in general, the set of products sold by the firm each period is sub-optimal in the sense that it does not maximise period profits (even though it is optimal in a dynamic perspective). In addition, in equilibrium, the existence of adjustment costs affects also the evolution of rival firms' product lines and of aggregate market variables and, *a priori*, the impact on firms' profitability is unclear. The results that I have obtained suggest however that the existence of adjustment costs does indeed decrease firms' profits.

The structural empirical models that I present in the first two Chapters of this thesis explain the range of qualities sold in the vertically-differentiated market for personal computers as the result of competition among a set of monopolistically competitive firms in a technologically dynamic environment. In these models, the range of qualities sold in the market is determined by competitive forces rather than by cannibalisation considerations.

However, cannibalisation may also be an important factor underlying firms' product line decisions in the presence of technological progress. In some cases, a firm may optimally add a new product to its product line, without withdrawing any lower-quality product ("product proliferation") thereby improving its ability to discriminate between consumers with different willingness to pay for quality. In some other cases, however, when a firm introduces a new product in its product line, it may decide to withdraw products of lower quality to avoid cannibalisation ("creative destruction").

In Chapter 3 I complement the analysis of the first two Chapters by presenting a vertical-differentiation model to study creative destruction as arising from optimal product line design by a monopolist. In particular, I consider a model in which firms can sell any of a given set of feasible qualities under general specifications of utility and technology. I characterise the impact that the introduction of a new quality in the product line has on other product-qualities and determine a necessary and sufficient condition for creative destruction (as opposed to product proliferation) to be optimal.

Inthis model, when creative destruction takes place because of price discrimination considerations it has a localised nature, i.e. the introduction of a product may drive products of neighbouring qualities out of the market.

This may well explain creative destruction in the form of product upgrading that occurs in those industries in which new products take the place of older similar versions (e.g. software, cars). However, it does not explain the process of creative destruction observed in the market for personal computers where movements of the frontier of technology are associated with the exit of "distant" products at the bottom of the quality ladder.

Creative destruction is a fundamental determinant of product lines in dynamic vertically-differentiated markets. The first three Chapters of this thesis study the economic mechanisms that determine if and to what extent products of lower quality survive the introduction of better products in the market. In a number of industries for durable goods, however, despite discontinuing the production of older-generation products, manufacturers often continue to (re)sell relatively obsolete products by acting as intermediaries in their exchange in the second-hand market. For instance, since 1987 IBM has been reselling used IT equipment, such as mainframes, midrange systems and PC servers. Other manufacturers in IT industries, such as Compaq, Dell, Hewlett Packard, Cisco and SUN offer trade-in promotions and resell used equipment, which they usually refurbish and often adapt to the specific needs of customers. In the market for aircraft, Boeing has been selling even used Airbuses. Most car manufacturers sell certified pre-owned vehicles to which they extend warranty coverage and now account for one third of used cars that are exchanged in the US.

In order to study the incentives of a durable-good manufacturer to trade used goods, in Chapter 4, I consider a market for a durable good subject to quality-improving technological change. Secondary trading has two main effects on the market for the new technology. Used goods are a substitute for new goods and hence tend to decrease the price at which the new technology can be sold. At the same time, however, second-hand trade makes it possible for former patrons to upgrade to the new technology (or in general to upgrade more frequently). I suggest that, in some industries, by intervening as an intermediary in the market for used goods and internalising the transaction costs of second-hand trade, the seller may ease the process that allows former patrons to upgrade to the new technology while selling the new technology to new consumers at the highest price possible and be better off than by leaving secondary trading to be managed by specialised intermediaries.

Chapter 1

Entry in differentiated markets: a structural model of optimal product lines

1.1 Introduction

Since Hotelling's (1929) seminal contribution, a large body of research has investigated product differentiation as the outcome of firms' optimal economic decisions.[1] However, despite the well-established theoretical literature, empirical structural studies of product differentiation have been few. In strategic settings, the potential multiplicity of equilibria and computational difficulties have been the main obstacles to the introduction of differentiated entry locations in empirical models of market structure. Some progress, however, has recently been made in the study of firms' entry decisions in differentiated mar-

[1] The theoretical literature on product differentiation is vast. Seminal contributions to this literature include Shaked and Sutton (1982, 1983 and 1987); Spence (1976a, 1976b) and Dixit and Stiglitz (1997).

kets as the equilibrium outcome of a discrete game played by a set of potential entrants (Mazzeo 2002, Seim 2005, Davis 2005).

In this Chapter, I develop and estimate a static structural model of endogenous quality choice, which I use to study the determinants of optimal product lines in the market for personal computers.

Like the existing empirical literature on entry, I develop a structural model that explains market structure as the outcome of firms' optimal (product) entry decisions, and I use the configuration of products that I observe in the market to draw inference on parameters of the profit function, and in particular to estimate structural parameters of firms' cost function.

However, the original modelling approach that I follow takes advantage of some features of the market for personal computers to simplify the analysis of firms' product-line decisions, and distinguishes this study from the existing empirical literature on endogenous market structure. First of all, since I deal with a rather fragmented market, I consider a model of endogenous quality choice under monopolistic competition rather than a strategic entry game. In addition, given the nature of product differentiation in the market for personal computers, and the availability of detailed data on product quality, I am able to treat firms' location decisions as a continuous variable rather than a discrete choice. Finally, I focus on the product location decisions taken by a number of established rivals, i.e. I do not model the number of firms active in the industry as endogenous.[2]

I use this structural model to study the range of product-qualities sold in the market for personal computers over a number of years, capturing the process, common to many innovative industries, that can be described as a

[2]However, the model could be easily extended to endogenise the number of active firms in the market.

"moving quality window": as technological progress allows firms to introduce products of better quality in the market, relatively obsolete products at the bottom of the quality ladder are withdrawn in a relatively ordered sequence.[3] The width of the quality window depends on the extent to which products of lower quality can profitably survive the introduction of better products of higher quality in the market, which in turn depends on the structural characteristics of demand and supply, such as consumers' preference for quality and production costs.

By developing and estimating a structural model, I can study how the range of product-qualities sold in the market is affected by changes to the characteristics of the market environment even in the absence of relevant natural experiments. In particular, I simulate the impact of a change in consumers' marginal utility of quality on the equilibrium range of product-qualities supplied in the market and I estimate the impact on firms' profits.

The development of software is an important determinant of the decline in the value of computers (Geske, Ramey and Shapiro 2004). A change in consumers' marginal utility of quality can be thought of as capturing the development of new software applications that increase the utility that consumers can derive from the quality, e.g. from the speed, of a PC. In this model, the impact of an increase in consumers' marginal utility of quality is to increase the sales of PCs of relatively high quality and to reduce the sales of PCs of relatively low quality. As a result, the equilibrium range of qualities supplied in the market is reduced because some low-quality products are drive out of the market and, in a dynamic perspective, the life cycle of a personal computer of any given quality is shortened.

I also study the impact that a reduction in the cost of providing product

[3]The term "moving quality window" was originally used by Sutton (2001).

variety in the form of a wide(r) range of qualities has on firms' profitability. In particular, I show that this change increases firms' profits, but that this increase is significantly reduced by firms' endogenous reaction, i.e. the increase in the equilibrium range of qualities sold in the market.

This Chapter is organised as follows. I first review the related literature. In section 1.3 I describe the industry. In section 1.4 I describe the model and its equilibrium and in section 1.5 I discuss the empirical specification and estimation of the model. Finally, in section 1.6 I present the results of the model simulations and final conclusions follow.

1.2 Related literature

Some recent contributions to the empirical literature on entry have extended the equilibrium models estimated by Bresnahan and Reiss (1990, 1991) and Berry (1992) to markets with non-homogeneous products.

Mazzeo (2002) considers a discrete entry model with complete information, which he uses to study entry decisions with endogenous quality choice in the oligopolistic motel markets along U.S. interstate highways. He shows that the negative effect that a competitor has on a firm's payoffs is up to twice as large if that competitor sells the same product type and provides empirical support to the product choice theories that predict firms will offer products unlike those of their competitors. However, a drawback to this model is that computing an equilibrium configuration is difficult for markets with large numbers of locations and firms: Mazzeo's analysis shows that even with only three quality levels, estimation becomes burdensome due to the large number of profit constraints that must hold in an equilibrium configuration.

Seim (2005) studies entry in differentiated markets with an incomplete-information model in which firms' have private shocks to their profitability that are not observed by competitors. This modelling approach allows her to derive the resulting Bayesian Nash equilibrium conjectures more easily than in Mazzeo's complete information framework and, as a result, it can potentially deal with a larger number of quality locations. She applies her model to the location decision of new firms in the video retail industry and obtains results that show evidence of strong incentives for spatial differentiation.

Davis (2005) focuses instead on the problem of the potential multiplicity of equilibria in strategic games of entry and provides both a generalisation and an extension of the theoretical results developed by Bresnahan and Reiss (1991) and Berry (1992). He considers a class of games in which firms' profit functions depend on their own strategy and their rivals' strategy only through an index of market output and derives sufficient conditions under which this index of market output is uniquely determined within the set of pure strategy market equilibria. This result can inform an empirical estimation strategy of this class of models even in the presence of multiple equilibria.

Like these studies of entry in differentiated markets, I also consider a model in which firms make endogenous product entry choices and I use observed market structure to draw inference on parameters of the profit function. However, there are two main differences between the approach that I follow and the existing literature: since I deal with a relatively fragmented market, I consider a model of monopolistic competition rather than a strategic game; and I treat a firm's location decision as a continuous variable rather than as a discrete variable. As a result, I am able to deal with (product) entry in a rather complex differentiated market with a large number of quality locations.

Since I focus on the product-quality decisions taken by a set of established rivals, this study can also be looked at as a structural analysis of optimal product-line decisions. The empirical literature on the determinants of optimal product-lines spans the boundaries between economics and marketing but is also rather limited, with the vast majority of existing research following a reduced-form approach.

Stavins (1995) investigates entry and exit decisions at the model level in the market for personal computers in the period between 1976 and 1988. In the marketing literature, Bayus and Putsis (1999) also provide a reduced-form analysis of additions to, and deletions from, product lines in the market for personal computers. More recently, Draganska and Jain (2005) introduce choices of product range (number of flavours) into a structural model of competition in the horizontally-differentiated market for yoghurt.

I contribute to this literature by developing an original structural model of quality choice that explains optimal product lines as the result of firms' optimal decisions. Empirical models of quality choice have also been developed by Crawford and Shum (2005), who study the welfare impact of endogenous quality choice in the market for cable television in the U.S using a modelling approach derived from the theoretical screening literature used in the analysis of optimal non-linear pricing, and by Carranza (2006), who develops a dynamic model of endogenous quality choice in the market for digital cameras which is somehow similar to the dynamic model that I will present in Chapter 2.

The market for personal computers has been the subject of various studies. In many the analysis has been focused on the estimation of the demand for personal computers. Genakos (2004) estimates a random coefficient discrete choice model of the demand for personal computers. Bajari, Benkard and Lanier (2005) estimate a hedonic model to estimate the demand for personal computers. Foncel and Ivaldi (2005) estimate a differentiated-products model of the home PC market. Stavins (1997) also employs hedonic coefficients to estimate price elasticities for differentiated products in the market for personal computers. Bresnahan, Scott and Trajtenberg (1997) evaluate the sources of temporary market power in personal computers in the late 1980s and find segmentation both between frontier and nonfrontier products and between branded and nonbranded products.

Recently, Deltas and Zacharias (2006) have also studied the pricing behaviour of PC manufacturers. Their results suggest that entry prices of firms are set to extract short run rents from consumers who have a higher willingness to pay for their brand. Stengos and Zacharias (2006) also study the intertemporal pricing of personal computers and find that high-quality firms charge higher premia only for their most advances products.

1.3 The personal computer industry

The personal computer industry is a prototypical example of an industry characterised by prominent vertical differentiation in which the range of the quality of the products sold in the market evolves over time as the result of a seemingly relentless process of "creative destruction" induced by ongoing technological progress.

In this paper, I use a dataset that comprises quarterly data on quantities, prices and characteristics of desktop personal computers sold by the top nine producers in the U.S. market in the home segment between 1995Q1 and 2001Q2. The data is from the personal computer tracker (PC Tracker), an industry census conducted by International Data Corporation (IDC).[4] IDC data on personal computers have been used in a number of studies on the PC industry, e.g. Foncel and Ivaldi (2005), Pakes (2003) and Bayus and Putsis (1999). The IDC dataset provides disaggregation by manufacturer, brand name, and

[4]The dataset was kindly made available to me by Christos Genakos and a detailed description of the constructed variables can be found in Genakos (2004).

processor froup (chip type and processor speed bandwidth). The processor groups are described in Table 1.1. Each observation in the IDC dataset (defined as a combination of vendor, brand and processor group[5]) was matched with more detailed product characteristics from various PC magazines. Because of the need to match the original dataset with additional information, the available data covers only the main 9 producers: Acer, Compaq, Dell, Gateway, IBM, Sony, Toshiba, Hewlett-Packard and NEC.[6] These firms account for about 60% of total desktop PC sales in the home-segment (see also Foncel and Ivaldi, 2005).

Table 1.3 provides a summary of prices, sales and quality of desktop personal computers over the 26 quarters covered by the dataset.

Prices of desktop personal computers have decreased by about 50% in real term in around 5 years. The price reduction is remarkable, especially in relation to the significant increase in the quality of personal computers over time. Figure 1-3 shows the evolution of prices by processor group and Table 1.5 shows the evolution of average prices by vendor.

Aggregate sales have followed an upward trend but also show a clear seasonal pattern as illustrated in Table 1.3. Figure 1-2 shows the evolution of aggregate sales by processor group, providing an overall picture of technological life cycles over the period considered. In the period covered by my dataset, important CPUs were introduced in the market (e.g. Pentium, Pentium II, Pentium III). Some of these also disappeared from the market before 2001Q2, e.g. Pentium, Pentium MMX, Pentium II. Therefore, the period covered by my dataset is sufficiently long to allow me to observe the entire life-cycle of

[5]Notice, therefore, that in general an observation in my dataset is the aggregation of different individual products.

[6]I do not have data for Sony and Toshiba for all periods. These firms, however, account only for 3% of the sales of the 9 manufactuers.

some important CPUs.

The market shares at the vendor level (calculated on the basis of their combined total sales) are reported in Table 1.4. Since these vendors account for about 60% of total desktop PC sales in the home segment, these figures actually overestimate the true market shares of the vendors.

The average market share of an observation in my sample (calculated relative to total sales) is equal to only around 1.5%. Moreover, 95% of the observations in my sample have a market share of less that 6.6%. These figures suggest that the average market share of a product in this market is very low since the statistics that I have reported overestimate the average market share of a product for two reasons: (i) each observation is actually the aggregation of a number of products and (ii) the vendors in my sample account only for about 60% of total desktop sales.

The quality of a personal computer derives from the combination of a number of characteristics such as the processor, the hard-disk, the RAM, etc. In Table 1.3, I capture quality by the speed of the PC processor (in MHz) which is a key determinant any personal computer's performance and is strongly correlated with other characteristics such as the amount of RAM and the size of the hard disk as shown by Table 1.2.[7]

The average speed of personal computers sold in the market has increased by about ten times between 1995 and 2001. This increase derives from both the introduction of computers of faster quality and the exit of computers at the bottom of the quality ladder from the market. As mentioned above, in the period covered by the dataset, new important generations of CPUs such as Intel Pentium, Intel Pentium II and Intel Pentium III were introduced in the

[7]The CPU industry has been the subject of a number of recent studies including Song (2007) and Gordon (2006).

market while older CPUs such as Intel 486, Intel Pentium and Intel Pentium II completed their life-cycles.

This process is summarised graphically in Figure 1-4, which shows the evolution of the range of qualities of PCs sold in the market in the period between 1995Q1 and 2001Q2. In the graph, quality is measured by the natural logarithm of the processor benchmark. The benchmark of a processor is a more encompassing measure of the processor's quality than speed because it accounts for differences in the processor architectures that affect performance beyond the simple measure of clock speed.

The graph shows that the evolution of the range of qualities sold in the market over time follows a process of "creative destruction" that can be described as a "moving quality window": as technological progress shifts the frontier of technology forward, obsolete products are progressively withdrawn from the market. The "moving quality widow" pattern seems common to many other innovative industries but not to all: for instance, de Figueiredo and Kyle (2006) show that the introduction of faster laser printers has not been associated with the exit of slower laser printers from the market.

Figure 1.9 shows in which of the exogenous processor groups defined in Table 1.1 each firm is active.[8] This analysis shows an important characteristic of the product lines sold by firms in this market: firms tend to sell product lines with no holes.[9] It is therefore reasonable to describe a firm's product line by the range of qualities it sells, rather than by the detailed specification

[8]Notice, therefore, that the quality variable in this graph is discrete by construction.

[9]Notice that most instances of "holes" are related to the Pentium Pro, which has always been a very marginal technology as can be seen from its limited sales in Figure 1-2. Towards the end of the sample, the holes in the product line are essentially related to the AMD Athlon processor which was just introduced in the market. These holes appear to be related more to the availability of AMD processors, which were not adopted by several manufacturers who chose to use Intel products, rather than to choices concerning the range of qualities sold by vendors.

of the quality of each product sold. Another feature of firms' product lines, which I will discuss in more detail in the next Chapter, is that firms do not appear to seek differentiation from the rivals in the quality space as there is no evidence of *persistent* inter-firm product differentiation over time.

1.4 The model

Demand is specified as a standard discrete-choice logit model. There are M consumers in the market. The utility that consumer i obtains from purchasing product j, $j = 1, 2, ..., n$, is specified as

$$u_{ij} = \alpha q_j + \gamma p_j + \varepsilon_{ij} = \delta_j + \varepsilon_{ij} \tag{1.1}$$

where q_j and p_j are respectively the quality and the price of product j, and ε_{ij} is an i.i.d. extreme-value distributed error term. It is reasonable to expect that $\alpha > 0$ and that $\gamma < 0$. Since individuals' heterogeneity enters the model only through the non-deterministic part of utility, utility can be decomposed into the mean-utility term $\delta_j \equiv \alpha q_j + \gamma p_j$ and the stochastic idiosyncratic component ε_{ij}.

Each consumer purchases one of the products sold in the market to maximise her utility. The market share of product j is given by the standard logit expression:

$$s_j = \frac{\exp(\delta_j)}{\sum_{i=1}^{n} \exp(\delta_i)} \tag{1.2}$$

There are n_f firms active in the market. The set of product-qualities that firms can introduce in their product lines is determined by exogenous technological conditions: there are N_j equally-spaced feasible qualities, spanning the range between 0 and the frontier of technology $\overline{Q}$. Let $\Delta q = 1/N_j$ be the width of each quality location. Each firm f sells a product line with no holes $\left[\underline{q}_f, \overline{q}_f\right]$ that includes all feasible qualities between the "marginal quality" $\underline{q}_f$ and the top quality $\overline{q}_f$.

A product of quality q can be produced at constant marginal cost

$$c(q) = \phi q$$

In addition, I assume that each firm incurs the fixed cost of production F, which does not depend on a firm's output but which depends on the range of product-qualities supplied:

$$F = F\left(\overline{q}_f - \underline{q}_f\right)$$

The profits of product-quality q_j sold by firm f are thus

$$\pi_f(q_j) = M\frac{\exp\left(\alpha q_j + \gamma p_f(q_j)\right)\left(p_f(q_j) - \phi q_j\right)}{S} \tag{1.3}$$

where $S \equiv \sum_{f=1}^{n_f}\sum_{q=\underline{q}_f}^{\overline{q}_f} \exp\left(\alpha q + \gamma p_f(q)\right)$ is the "aggregate index of market output", which fully summarises the impact that market structure has on the market share, and hence on the profitability, of any product-quality.

Each firm sequentially chooses (i) its product line $\left(\underline{q}_f, \overline{q}_f\right)$ and (ii) the price of each product included in its product line. I analyse the equilibrium of this model under the following set of assumptions.

Assumption 1.1 *A product's profitability is increasing in quality:* $(\alpha + \gamma\phi) > 0$

Assumption 1.2 *The introduction and the pricing of any product has a* *negligible impact on the aggregate index of market output.*

Assumption 1.3 *There is a continuum of feasible quality locations:* $\Delta q \rightarrow 0$.

Assumption 1.1 imposes a restriction on the parameter space that ensures that a product's profits are monotonically increasing in quality and therefore that, in equilibrium, optimal product lines include only products of relatively high quality. This assumption can be verified empirically given estimates of α, γ and ϕ.

Assumption 1.2 entails the consideration of a monopolistically competitive market in which each firm's decisions have a negligible impact on aggregate market variables. This assumption implies that (i) firms do not interact strategically in the market when taking product-line and pricing decisions and (ii) cannibalisation considerations do not play a role in firms' pricing and product-line decisions. This is because in this model all cannibalisation effects are captured by the impact that (the price of) a product has on the aggregate index of market output.[10] The assumption of monopolistic competition is potentially restrictive but it is consistent with the substantial fragmentation at the product level in the market for personal computers and with assumption 1.3.

Assumption 1.3 states that there is a continuum of feasible qualities. This assumption is important because it allows me to treat quality as a continuous variable when I analyse firms' product-line choices.[11] Notice that, since firms sell all feasible qualities in an interval, this assumption also implies that each firm is selling a continuum of product-qualities. In the limiting case in which each firms sells a continuum of qualities, firm f's profit function can be written

[10] In Chapter 3, I consider a vertically-differentiated model to focus on the analysis of cannibalisation considerations and product-line decisions in the presence of product innovation.

[11] Mazzeo (2002) considers instead only three available quality locations.

as

$$\pi_f = \frac{M}{S} \int_{\underline{q}_f}^{\overline{q}_f} \exp\left(\alpha q + \gamma p\left(q\right)\right)\left(p\left(q\right) - \phi q\right) dq - F\left(\overline{q}_f - \underline{q}_f\right) \tag{1.4}$$

where

$$S \equiv \sum_{f=1}^{n_f} \int_{\underline{q}_f}^{\overline{q}_f} \exp\left(\alpha q + \gamma p\left(q\right)\right) dq$$

is the aggregate index of market output in the case in which each firm f sells a continuum of qualities between $\underline{q}_f$ and $\overline{q}_f$. The derivation of expression 1.4 can be found in Appendix 1.10.1.

The equilibrium of this model is a set of prices and of product lines such that (i) each firm maximises profits taking the aggregate index of market output as given and (ii) the aggregate index of market output is consistent with firms' optimal choices. The following proposition characterises equilibrium prices and product lines.

Proposition 1.1 *(i) In equilibrium, a product of quality q_j is sold at price equal to*

$$p\left(q_j\right)^* = \phi q_j - \frac{1}{\gamma} \tag{1.5}$$

(ii) Each firm sells all qualities in the range $\left[\underline{q}^*, \overline{Q}\right]$ where $\underline{q}^*$ solves

$$\underline{q}^* = \frac{1 + \ln\left[-\frac{\gamma F}{M}\left(n_f \int_{\underline{q}^*}^{\overline{Q}} \exp\left(\alpha q - 1 + \gamma \phi q\right) dq\right)\right]}{a + \gamma \phi} \tag{1.6}$$

Proof. See Appendix 1.10.2. ■

Notice that, given the assumption of monopolistic competition, optimal pricing entails a fixed markup above the marginal cost of production. The optimal markup does not depend upon the number of products or, more generally, on the set of product-qualities sold in the market and it does not converge to zero as the number of products in the market increases.[12]

Since I have not considered firm-specific product attributes or differences in marginal costs across firms, the equilibrium marginal quality $\underline{q}^*$ does not vary across firms. The model could be easily extended to deal with heterogeneous firms by introducing a "brand" characteristic or firm-specific marginal costs of quality. However, I have not found that in the market for personal computers there is any persistent difference across firms in the product lines supplied and I have therefore chosen to focus on the simplest specification of the model with no firm heterogeneity.[13]

Finally, since a product's profits are increasing in quality, a firm will sell all feasible qualities above the marginal one, up to the exogenous frontier of technology $\overline{Q}$.

1.5 Empirical specification and estimation

Estimation of the structure of model entails recovering the parameters of the utility function α and γ, the marginal cost of quality ϕ and the fixed cost F. I adopt a multi-step estimation strategy. First of all, I recover the parameters of the utility function by estimating a standard logit demand system. Then, by using the first-order-condition underpinning optimal pricing decisions, I estimate the marginal cost of quality. Finally, I infer the fixed cost F by matching observed and predicted optimal product lines, and in particular observed and predicted marginal product-qualities $\underline{q}$.

[12] This is consistent with Bajari and Benkard (2003) who observe that in a logit model with a large number of products Bertrand-Nash markups do not converge to zero.

[13] I discuss inter-firm heterogeneity further in Chapter 2

This sequential estimation approach is commonly used and it has both advantages and disadvantages compared to the simultaneous estimation of all the parameters of the model.

Since it does not enforce the equilibrium conditions that derive from the supply-side of the model in the estimation of demand, consistency of the estimates of the demand parameters does not depend on supply-side equilibrium assumptions and is therefore robust to a wide set of possible assumptions.

However, conditional on the model reflecting the true data-generating process, greater efficiency could be obtained by imposing the supply-side equilibrium conditions in the estimation of demand parameters. In addition, using the full structure of the model could also address some potential endogeneity problems in the estimation of demand. For instance, estimating demand on the basis of a selected panel sample in which attrition is driven by endogenous product-line choices without accounting for the endogenous selection results in inconsistent and biased estimates. Using an explicit model of product entry and product exit to explain this endogenous attrition could potentially overcome this problem. This is potentially an important benefit of developing a model of endogenous product choice and I plan to explore this issue in future research.

1.5.1 Estimation of demand

In the model, demand is specified as a standard discrete-choice logit model. The estimation equation that I use estimate the parameters α and γ is based on expression 1.2 and is derived (in Appendix 1.10.3 as the difference between the logarithm of the market share of a given product ($\ln(sh_{jt})$) and the average of the logarithms of all products' market shares $\left(\overline{\ln (sh)_t}\right)$:

$$\ln (sh_{jt}) - \overline{\ln (sh)_t} = \alpha (q_j - \overline{q}) + \gamma (p_j - \overline{p}) + \epsilon_{jt} \tag{1.7}$$

where market share is defined as sales of product j over total observed sales and ϵ_{jt} is an i.i.d. error term.[14]

I define the quality of good q_{jt} as the natural logarithm of the processor benchmark of the product.[15] Processor speed is the most important characteristic of a personal computer and it is strongly correlated with other characteristics, such as the amount of RAM and the size of the hard disk. Benchmarks are essentially numbers assigned to each processor-speed combination based on technical and performance characteristics. They are very strongly correlated with the speed of the processor but capture also other characteristics of a processor that affect performance and are therefore a more encompassing measure of quality.[16]

To the extent that there are unobserved factors affecting demand which are not observed by the econometrician but which are observed by firms in advance of setting prices, prices will be endogenous and estimation methods based on ordinary least squares (OLS) will not generate unbiased and consistent estimates.

For instance, endogeneity may arise because there are demand shocks unobserved by the econometrician but observed by the firm. I follow the approach that was originally proposed by Hausman, Leonard and Zona (1994)

[14]Notice that the estimates obtained from this equation do not differ from those obtained by the more common estimating equation for the logit demand model defined on the basis of the difference between the market share of a product and the market share of the outside option and including a full set of time dummies.

[15]CPU benchmarks were obtained from The CPU Scorecard (www.cpuscorecard.com).

[16]CPU benchmarks have been widely used as a measure of quality. See for instance Bajari and Benkard (2004) and Genakos (2005).

and Hausman (1996) which relies on using the prices of goods in other markets as instruments for the price of goods in a given market. Specifically, I use the Canadian price of a product as an instrument for its US price.

The underlying identifying assumption is that demand shocks are not correlated across markets while cost shocks are. If this is true, prices in another market therefore should be correlated with the price in the given market but should not be correlated to demand. This class of instruments has been used commonly used in the recent IO literature on the estimation of demand for differentiated products, not only in product-level models of demand but also in discrete-choice characteristics-based models of demand, e.g. Nevo (2000).

This instrument, however, is not suitable to correct for the potential endogeneity that arises if prices are correlated with unobserved product characteristics, since these unobserved characteristics are likely to be the same in different markets. This could potentially be a problem given that I do not control for many product attributes. However, the measure of quality that I use is strongly correlated with other important product characteristics that I do not include as regressors in the estimating equation, such as the size of the hard-disk and RAM. I also consider two alternative specifications of the model in which I control for as many product characteristics as possible to assess if, despite this strong correlation, unobserved product characteristics still affect the estimate of γ obtained from equation 1.7.

Other instruments have also been in the literature to address the endogeneity of prices. In particular, instruments can also be derived by considering variables that are correlated with the price-cost margin (rather than the cost). For instance, if product characteristics and product lines are considered exogenous, instruments could be constructed on the basis of the characteristics of other products sold in the market and of the ownership structure of products (see Berry, Levinsohn and Pakes, 1995). I do not attempt to use the qualities of other products as an instrument because in this model I do not consider product qualities to be exogenous.

The estimation results are presented in Table 1.6 where the IV specification that I use is reported in the second column. The Table reports also OLS estimates. The IV price coefficient estimate is consistent with the results obtained in other studies on the demand for personal computers that have used a similar demand model (see for instance the logit estimates in Genakos (2004)). The result of the Durbin-Wu-Hausman test for endogeneity: suggests that it is possible to reject the hypothesis that OLS estimates are consistent at the 1% level and therefore provides support to the use of instruments.

In order to assess if unobserved product characteristics are the source of further endogeneity problems, I have also run the estimation with product fixed-effects, which are meant to control for all unobserved product characteristics that are time invariant, and without product fixed-effects but controlling for all product characteristics that I observe in my dataset.[17] The estimates of γ obtained in these two models, which are also reported in Table 1.6, are not materially different to the IV estimate derived in the simpler regression in which I use one product characteristic only to control for quality.

The analytical simplicity of the logit specification of demand clearly comes at the price of substantial and possibly undesirable structure being imposed on consumers' substitution patterns. It is well known that the logit model assumes that the distribution of a consumer's preferences over products other than the product bought do not depend on the characteristics of the preferred product and therefore that own price derivatives depend only on market shares.

[17] The additional product characteristics include: RAM, hard disk size, presence of a CD-ROM, speed of modem, monitor, monitor size, presence of a ethernet card, dvd reader and soundcard.

Despite its limitations, however, the logit demand specification is still commonly used in empirical structural models because of its analytical tractability. In particular, the analysis of optimal product lines in this Chapter and Chapter 2 is greatly simplified by the property of the logit demand function whereby a product's demand depends on the other products sold in the market only through an aggregate index of market output, a property that neither the random-coefficient logit (Berry, Levinsohn and Pakes (1995)) nor the pure-characteristic model of demand (Berry and Pakes (2005)) share.

1.5.2 Estimation of the marginal cost of quality

I have specified a marginal cost function with constant returns to quality

$$c_t(q_i) = \phi_t q_i$$

Since firms buy PC components in an open market, the cost of producing a given product-quality may also be changing exogenously over time and I allow the marginal cost of quality to be time-dependent.

The estimating equation is derived from the first-order-condition for optimal pricing, i.e. from expression 1.5. More specifically, given the (IV) estimated marginal utility of income $\widehat{\gamma} = 2.414$ and observed prices, I run the following hedonic-type regression:

$$p(q_i) + \frac{1}{\widehat{\gamma}} = \phi_t q_i + u_{it}$$

where u_{it} is an error term assumed to be independent and identically distributed across qualities and time and drawn from a standard Normal distribution. The regression has no intercept and time-dependent slopes that capture the variation of the marginal cost of quality over time. Table 1.7 shows the results of the estimation.

The identification of the marginal costs of production is based on the optimal price-cost margins predicted by the theoretical model, and therefore depends on the assumption that consumers' marginal utility of income is constant and on the assumption that firms are monopolistically competitive, which implies an optimal constant markup over the marginal cost of production.

1.5.3 Estimation of the fixed cost

The last parameter that I need to estimate is the fixed cost of product variety F, which I identify on the basis of the marginal qualities $\underline{q}$ sold by firms. Proposition 1.1 states that the marginal quality sold by firms is

$$\underline{q}^* = \frac{1 + \ln\left(-\frac{\gamma F S^*}{M}\right)}{\alpha + \gamma\phi} \tag{1.8}$$

I use expression 1.8 to construct a simple non-linear OLS regression to estimate the parameter F as follows.

Given estimates of α, γ and of ϕ and firms' observed product lines, I construct an estimate of the index of market activity S^* as

$$\widehat{S}^* \equiv \sum_{f=1}^{n_f} \int_{\underline{q}_f}^{\overline{q}_f} \exp\left(\widehat{\alpha} q - 1 + \widehat{\gamma}\widehat{\phi} q\right) dq$$

I set "market size" M equal to the number of total sales each period. Estimation results are reported in Table 1.8.[18] Given that the average length of firms' product lines in the sample is 0.88, the estimated average fixed cost of production is therefore estimated to be about \$44 millions per quarter. This

[18]Notice, however, that the standard error is not corrected for the sample bias that affects the parameters α, γ and ϕ.

fixed cost is about 41% of average gross period profits predicted by the model.

1.6 Simulation of equilibrium product lines

Given the estimates of the parameters of the model that I have recovered empirically, I can simulate the equilibrium range of qualities sold in the market. In particular, by considering exogenous changes to the technology, i.e. different values of $\overline{Q}$ and of ϕ, and to market size, I can simulate the evolution of the equilibrium range of qualities over time so as to reproduce the moving quality window that I observe in the market in the period between 1995 and 2001.

In order to determine the equilibrium range of qualities supplied in the market, I solve numerically the following expression for $\underline{q}^*$ for each period, i.e. for each vector of exogenous variables/parameters $\left[\overline{Q}, M, \phi\right]$:

$$\underline{q}^* = \frac{1 + \ln\left[-\frac{\gamma F}{M}\left(n_f \int_{\underline{q}^*}^{\overline{Q}} \exp\left(\alpha q - 1 + \gamma\phi q\right) dq\right)\right]}{a + \gamma\phi} \tag{1.9}$$

In the simulations, I set M equal to observed total sales in each period, and I consider the frontier of technology to be equal to the average top quality sold by firms in the market.

The actual and predicted marginal qualities are shown in Figure 1-5, together with the estimated exogenous frontier of technology. It can be seen that, despite its simplicity, the model tracks reasonably well the evolution of the range of qualities sold in the market, even though the simulated marginal quality shows a seasonal behaviour, which is due to the seasonality in the total sales that I observe in the market.

One of the main advantages of estimating a structural model is the ability to simulate counterfactual scenarios to study economic phenomena in the absence of relevant natural experiments. In particular, I use the model to study the impact of a change in consumers' willingness to pay for quality. A change in consumers' marginal utility of quality can be thought of as capturing the development of new software applications that increase the utility that consumers can derive from the quality, e.g. from the speed, of a PC. As I show in Appendix 1.10.4, in this model, the impact of an increase in consumers' marginal utility of quality is to increase sales of personal computers of relatively high quality and to reduce sales of personal computers of relatively low quality.

The simulated impact of a an increase in α from the estimated value of 2.06 to the value of 2.5 on the range of qualities sold in the market is shown in Figure 1-6.[19] Figure 1-6 shows that, as expected, the increase in consumers' willingness to pay for quality increases the level of the marginal quality in the market, i.e. reduces the width of the quality window. I have estimated that the impact of the increase in consumers' willingness to pay on a firm's profits is on average \$9.8 millions per quarter, relative to an estimated average quarterly net status-quo profits of about \$61 millions.

I have also simulated the impact of a change in F, i.e. the cost of providing product variety in the form of a wider range of qualities. A change in F has both a direct and an indirect effect on a firm's profitability. First of all, for any given range of qualities sold, a change in F has a clear impact on the fixed cost borne by a firm. However, a change in F also affects the product line sold by a firm and by its rivals in equilibrium, which could partially offset the direct impact on profitability. The simulated impact of a 50% decrease in

[19] Notice, however, that this counterfactual analysis is undertaken for a given level of total sales. Since in this model no consumer purchases the outside option in equilibrium when there is a continuum of products in the market, the counterfactual analysis does not capture the impact that a change in the willingness to pay for quality has on total sales.

F is shown in Figure 1-7. As expected, it can be seen that the equilibrium width of the "quality window" increases as it is now less expensive to sell any given range of qualities in the market. If the change in F did not affect firms product lines, a firm's profits would increase by $22.6 millions per quarter on average. However, taking into account the indirect effect, i.e. the endogenous changes to firms' product lines, the increase in profits due to the reduction in F considered is 39% lower, being only $13.8 millions per quarter.

Ideally, I would also like to use the model to study and to measure the impact of changes in the economic environment on consumer welfare, accounting for the endogenous determination of the equilibrium range of qualities sold in the market. However, the logit model belongs, together with the random coefficient logit discussed in Berry, Levinsohn and Pakes (1995) and the random coefficient probit (Hausman and Wise 1978, McFadden 1981), to a class of discrete-choice models with an error additive component with full support on the real line that are not adequate to quantify welfare in markets with a large number of products. This is because in these model the error term captures idiosyncratic tastes for products that are independent of product characteristics. Since the idiosyncratic shock has unbounded support, when there is a continuum of product-qualities available for purchase, the maximum utility that a consumer can expect to obtain in the market is infinite.

This limitation of the model could be overcome by the use of an alternative demand model that is better suited to deal with markets in which there is a large number of products. One potential alternative to the logit is the pure-characteristic demand model developed by Berry and Pakes (2001) and by Bajari and Benkard (2005). However, using a demand system of this kind would substantially complicate the derivation of optimal product lines because of the localised nature of consumers' substitution patterns captured by these demand systems, which does not make possible to use a simple aggregate index of market output to summarise market structure. Another possible alternative could be the use of a quality-adjusted representative-consumer demand model à la Dixit and Stiglitz (1977), which is often used in theoretical models of monopolistic competition with a continuum of products but which has not been widely used in empirical applications in industrial organization.

1.7 Conclusions

In this Chapter, I have presented an empirical structural model of endogenous product choice in a differentiated market. By considering a model of monopolistic competition with a continuum of feasible qualities rather than a strategic discrete entry game, I have been able to develop a manageable structural model of quality choice that captures the main economic determinants of optimal product lines in the complex and dynamic market for personal computers.

Despite its relative simplicity, the model reproduces relatively well the evolution of the aggregate range of qualities over time observed in the market for personal computers in the form of a moving "quality window", capturing the process of creative destruction whereby technological progress that allows firms to introduce new better products in the market also forces the exit of low-quality products.

I have recovered the structure of the model empirically. As in other studies of endogenous market structure, I have used the observed configuration of products sold in the market to draw inference on parameters of the profit function, and in particular on firms' production costs.

The model tracks reasonably well the evolution of the range of product-qualities sold in the market for personal computers. I have also simulated the impact of an increase in consumers' willingness to pay for quality and of a reduction in the cost of providing product variety on equilibrium product lines and firms' profitability. The impact of a change of consumers' willingness to pay for quality is to increase sales of relatively high-quality products and to reduce the sales of relatively low-quality products. As a result, when consumers value quality more, the range of products that can survive in the market is reduced, which also implies that product life-cycles are shortened. The impact of a reduction in the cost of providing product variety is to increase profits, but firms' endogenous modifications to their optimal product lines reduces the profit increase by 39% relative to the case in which product lines did not change.

The model can be developed and improved in a number of ways. For instance, the model does not endogenously explain technological progress, i.e. the shifts of the frontier of technology over time. Considering technological progress as exogenous is not an unreasonable assumption since in the market for personal computers much of the improvements to product quality stem from innovation that takes place in the upstream components markets. By endogenising technological progress, however, the model could potentially to be applied to a larger variety of markets in which innovation is an important feature of competition and to explain endogenous product variety in a more encompassing way. However, the only meaningful way to capture endogenous technological progress is in the context of a dynamic model.

Another reason why it would be useful to consider a dynamic model of optimal product lines is that introducing products to and withdrawing products from a firm's product line is normally not costless. In the next Chapter, I discuss how the consideration of product-line adjustment costs can improve the understanding of the range of qualities optimally sold by firms in the market for personal computers and I develop a dynamic model of optimal product lines that accounts for these adjustment costs.

1.8 Tables

Table 1.1: Processors and CPUs

Code	Processor group	CPU
40	486	Intel 486DX2, AMD 486DX2
50	5th generation $\leq$100MHz	Pentium
51	5th generation 101-149 MHz	Pentium, Pentium MMX, Cyrix Media GX
52	5th generation 150-179 MHz	Pentium, Pentium MMX, AMD K6
53	5th generation $\geq$ 180 MHz	Pentium, Pentium MMX, Cyrix Media GX
61	6th generation $\leq$ 200 MHz	Pentium Pro
62	6th generation 201-299 MHz	Pentium II, AMD K6
63	6th generation 300-399 MHz	Pentium II, Celeron, AMD K6
64	6th generation 400-499 MHz	Pentium III, Pentium II, Celeron, AMD K6
65	6th generation 500-599 MHz	Pentium III, Celeron, AMD K6-2
66	6th generation 600-699 MHz	Pentium III
67	6th generation 700-799 MHz	Pentium III, Celeron
68	6th generation 800-899 MHz	Pentium III, Celeron, AMD K6
69	6th generation 800-899 MHz	Pentium III
70	6th generation $\geq$ 1 GHz	Pentium III
76	7th generazion 600-699 MHz	Athlon
77	7th generazion 700-799 MHz	Athlon
78	7th generation 800-899 MHz	Athlon
79	7th generazion 1-1.5 GHz	Athlon, Pentium IV

Table 1.2: Correlation between main product characteristics

	speed	ram	hard
speed	1		
ram	0.87	1	
hard	0.88	0.79	1

Table 1.3: Quality, prices and sales in the market for Personal Computers

Period	Lowest speed	Average speed	Top speed	Lowest price	Average price	Top price	Total sales
1	50	85	133	1256	2333	4093	995179
2	50	89	133	1168	2205	3256	807281
3	50	93	133	1080	2161	3239	1219146
4	50	101	180	1359	2177	3897	1536614
5	50	123	200	1063	2227	4930	1219411
6	66	128	200	1250	2302	3903	960013
7	66	141	200	1338	2297	3632	1417532
8	90	155	200	1039	2083	3515	1723069
9	120	172	233	1046	2093	3773	1463808
10	120	191	266	984	2071	3562	1241581
11	120	196	266	884	1906	3198	1694879
12	166	229	350	929	1973	3346	2052952
13	150	270	450	831	1996	3138	1672734
14	166	292	450	944	1836	2807	1358593
15	180	306	450	854	1683	2515	1810164
16	180	323	450	765	1517	2497	2346936
17	233	395	550	867	1527	2347	2147084
18	266	418	550	764	1417	2223	1843719
19	300	428	550	672	1263	2239	2300758
20	300	489	650	623	1291	2270	2988577
21	400	616	1100	707	1422	2563	2904770
22	400	650	1100	658	1373	2237	2118870
23	400	665	1100	602	1312	2329	2694950
24	400	690	1100	521	1242	2267	2855113
25	400	727	1100	499	1177	2172	2025285
26	500	862	1500	534	1171	1949	1499484

Table 1.4: Vendors' market shares

	Acer	Compaq	Dell	Gateway	Hewlett-Packard	IBM	NEC	Sony	Toshiba
1995Q1	4%	12%	2%	16%	0%	12%	53%		
1995Q2	6%	15%	2%	14%	1%	13%	48%		
1995Q3	6%	12%	2%	11%	6%	13%	50%		
1995Q4	11%	13%	2%	11%	6%	11%	46%		
1996Q1	7%	17%	3%	12%	7%	7%	46%		
1996Q2	8%	15%	3%	15%	5%	12%	43%		
1996Q3	5%	23%	3%	10%	6%	9%	36%	4%	5%
1996Q4	5%	25%	4%	11%	5%	11%	33%	2%	4%
1997Q1	6%	20%	3%	15%	6%	6%	35%	3%	5%
1997Q2	5%	16%	4%	14%	8%	8%	39%	4%	3%
1997Q3	4%	30%	3%	15%	9%	6%	29%	3%	2%
1997Q4	6%	25%	4%	19%	8%	10%	25%	3%	1%
1998Q1	6%	22%	5%	19%	15%	8%	22%	3%	
1998Q2	2%	19%	7%	17%	16%	8%	28%	3%	
1998Q3	2%	20%	6%	20%	18%	11%	22%	1%	0%
1998Q4	3%	22%	5%	24%	17%	10%	17%	1%	0%
1999Q1	1%	23%	6%	24%	23%	8%	14%	1%	0%
1999Q2	1%	26%	9%	21%	19%	9%	14%	1%	0%
1999Q3	1%	27%	9%	23%	23%	7%	10%	1%	0%
1999Q4	0%	30%	11%	23%	27%	3%	4%	0%	1%
2000Q1	1%	32%	8%	24%	31%	1%	3%	0%	0%
2000Q2	1%	26%	11%	25%	33%	2%	1%	1%	0%
2000Q3	1%	32%	9%	24%	32%	1%	0%	1%	0%
2000Q4	1%	28%	11%	23%	35%	1%	0%	1%	0%
2001Q1	1%	24%	14%	27%	32%	0%	0%	1%	0%
2001Q2	1%	21%	20%	22%	32%	1%	0%	3%	0%
Mean	4%	22%	6%	18%	16%	7%	24%	2%	1%

Table 1.5: Average prices by vendor

	Acer	Compaq	Dell	Gateway	Hewlett-Packard	IBM	NEC	Sony	Toshiba
1995Q1	2,710	2,234	2,398	2,524	3,224	2,064	1,894		
1995Q2	2,693	2,094	2,398	2,313	2,725	1,899	1,826		
1995Q3	2,504	1,985	2,417	2,216	2,566	1,974	1,774		
1995Q4	2,012	2,283	2,350	2,443	2,459	2,043	1,829		
1996Q1	2,029	2,044	3,117	2,174	2,266	2,348	2,037		
1996Q2	1,762	1,979	2,562	2,709	2,292	2,236	2,071		
1996Q3	1,626	1,997	2,663	2,545	2,270	2,487	2,044	2,678	2,918
1996Q4	1,769	2,096	2,370	2,413	2,015	2,161	1,788	2,280	1,538
1997Q1	1,931	2,002	2,054	2,350	2,051	1,936	2,172	1,803	1,691
1997Q2	1,976	2,068	2,108	2,271	1,926	1,925	2,184	1,617	1,504
1997Q3	1,647	1,868	1,893	2,092	1,788	1,655	1,961	1,947	1,691
1997Q4	1,760	1,925	2,032	2,124	1,958	1,841	2,018	1,932	1,527
1998Q1	1,846	1,917	2,123	2,099	1,813	2,095	2,021	1,979	
1998Q2	1,721	1,672	1,899	1,905	1,642	2,044	1,941	1,959	
1998Q3	1,540	1,514	1,780	1,823	1,466	1,832	1,740	1,755	1,740
1998Q4	1,341	1,450	1,732	1,714	1,289	1,723	1,519	1,365	1,722
1999Q1	1,478	1,394	1,699	1,714	1,392	1,660	1,297	1,607	1,630
1999Q2	1,209	1,272	1,548	1,631	1,337	1,317	1,319	1,439	1,494
1999Q3	1,095	1,171	1,390	1,347	1,164	1,133	1,335	1,291	1,341
1999Q4	889	1,252	1,365	1,375	1,181	1,093	1,363	1,399	1,365
2000Q1	1,242	1,251	1,361	1,467	1,272	1,364	1,588	1,687	1,436
2000Q2	973	1,187	1,408	1,364	1,123	1,306	1,538	1,722	1,335
2000Q3	903	1,187	1,357	1,354	1,034	1,233	1,374	1,557	1,504
2000Q4	861	1,147	1,267	1,293	968	1,129	1,255	1,483	1,426
2001Q1	812	1,052	1,319	1,198	931	1,094	1,288	1,420	1,342
2001Q2	819	1,075	1,299	1,229	1,062	960	1,143	1,358	1,227
Mean	1,583	1,658	1,920	1,911	1,739	1,714	1,705	1,714	1,580

Table 1.6: Estimation of demand

Variable	OLS Estimates	IV Estimates	IV Fixed-effect Estimates	IV Estimates with all characteristics
Price (γ)	-1.807^{***}	-2.414^{***}	-2.531^{***}	-2.23^{***}
	(0.19)	(0.21)	(0.56)	(0.21)
Quality (α)	1.469^{***}	2.063^{***}		1.87^{***}
	(0.19)	(0.25)		(0.24)
RAM				-0.01^{***}
				(0.00)
Hard disk				0.01^{*}
				(0.01)
CD-ROM				0.47^{**}
				(0.24)
Modem				0.01^{***}
				(0.00)
Ethernet				-1.05^{***}
				(0.15)
DVD				1.34^{***}
				(0.15)
Sound card				-0.20
				(0.14)
Monitor				1.75^{***}
				(0.39)
Monitor size				-0.09^{***}
				(0.02)
R^2	0.11	0.16	0.1	0.29
N	1,703	1,703	1,703	1,703

standard errors in brackets

* : significant at the 10% level

** : significant at the 5% level

*** : significant at the 1% level

Table 1.7: Estimation of marginal costs

Parameter	Estimate	Standard error	Confidence interval	
ϕ_1	384.5	13.5	358.0	411.0
ϕ_2	354.2	12.3	330.1	378.4
ϕ_3	341.8	12.1	318.1	365.4
ϕ_4	340.9	11.5	318.3	363.4
ϕ_5	337.4	12.5	312.8	361.9
ϕ_6	344.1	12.2	320.2	368.1
ϕ_7	335.3	11.4	312.9	357.8
ϕ_8	293.8	10.3	273.5	314.1
ϕ_9	284.2	10.1	264.4	304.0
ϕ_{10}	275.3	8.5	258.6	292.0
ϕ_{11}	246.3	8.3	230.0	262.6
ϕ_{12}	248.6	7.8	233.2	264.0
ϕ_{13}	246.7	7.6	231.7	261.7
ϕ_{14}	217.8	7.9	202.3	233.4
ϕ_{15}	192.7	8.1	176.7	208.7
ϕ_{16}	164.9	8.5	148.2	181.5
ϕ_{17}	160.4	8.7	143.4	177.5
ϕ_{18}	143.0	9.1	125.2	160.8
ϕ_{19}	120.5	9.2	102.4	138.6
ϕ_{20}	121.4	8.7	104.3	138.5
ϕ_{21}	134.7	6.4	122.2	147.2
ϕ_{22}	126.9	6.6	114.0	139.9
ϕ_{23}	118.6	6.6	105.6	131.5
ϕ_{24}	108.9	6.2	96.7	121.0
ϕ_{25}	99.4	6.8	86.2	112.7
ϕ_{26}	96.3	6.1	84.4	108.3
R^2=0.94				
N=1,703				

Table 1.8: Estimation of fixed cost

Parameter	Estimate	Standard error	Confidence interval	
F	49,924	2,051	45,887	53,969
N=182				

1.9 Figures

Figure 1-1: Analysis firms' product lines

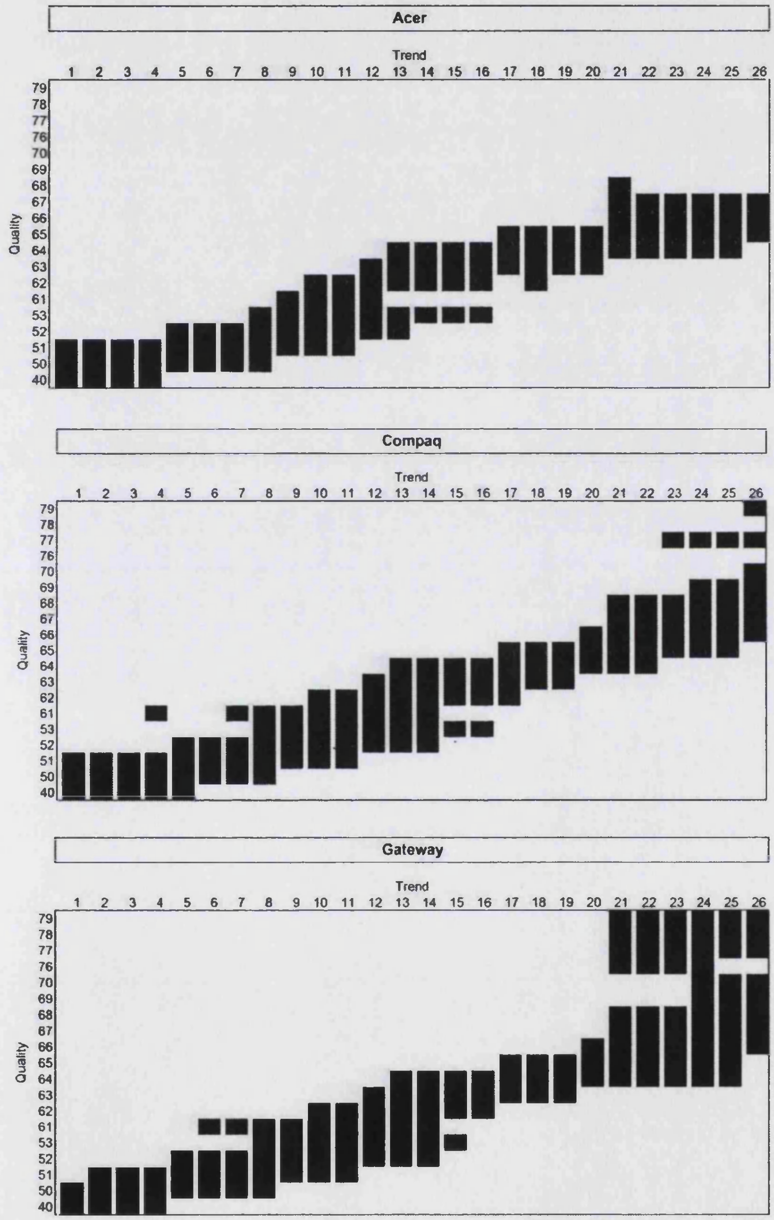

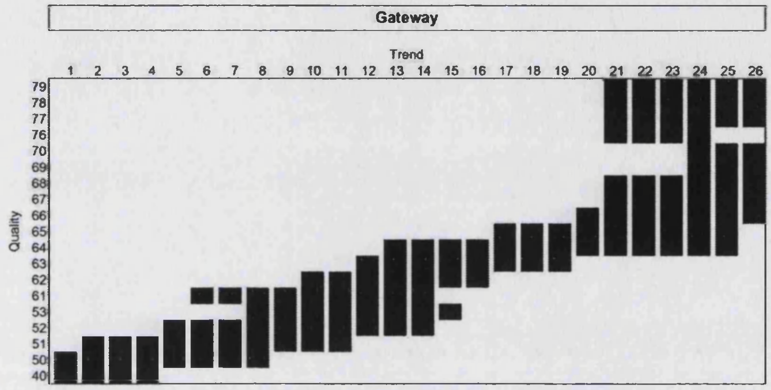
Gateway
Trend
1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26
79 78 77 76 70 69 68 67 66 65 64 63 62 61 53 52 51 50 40
Quality

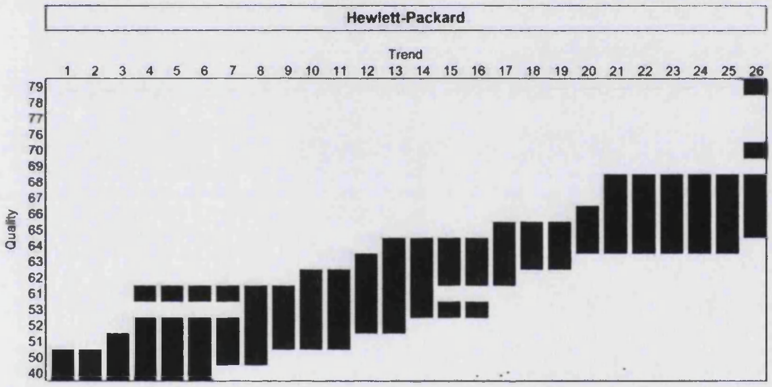
Hewlett-Packard
Trend
1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26
79 78 77 76 70 69 68 67 66 65 64 63 62 61 53 52 51 50 40
Quality

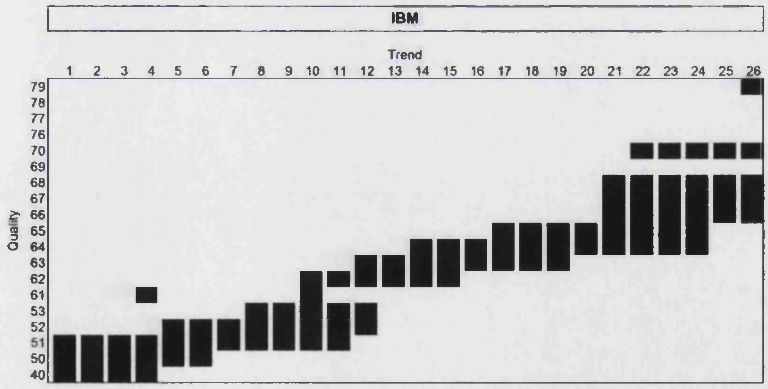
IBM
Trend
1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26
79 78 77 76 70 69 68 67 66 65 64 63 62 61 53 52 51 50 40
Quality

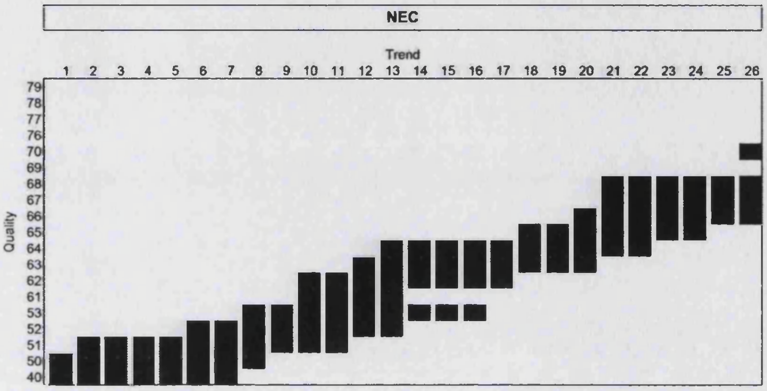
NEC
Trend
Quality

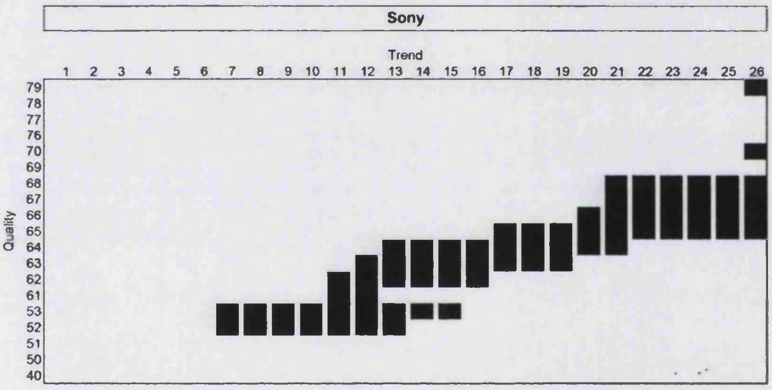
Sony
Trend
Quality

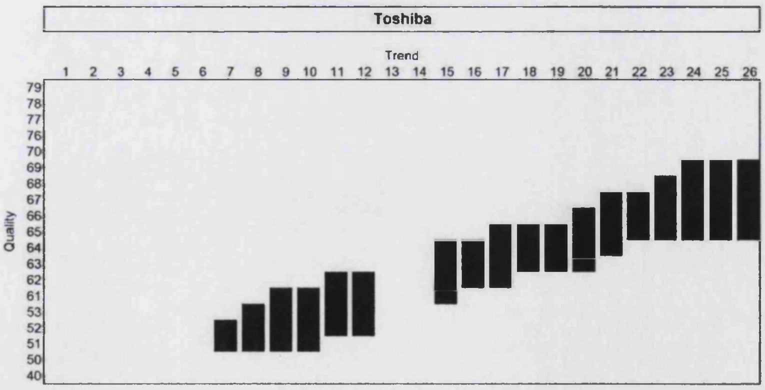
Toshiba
Trend
Quality

Figure 1-2: Aggregate sales by processor generation

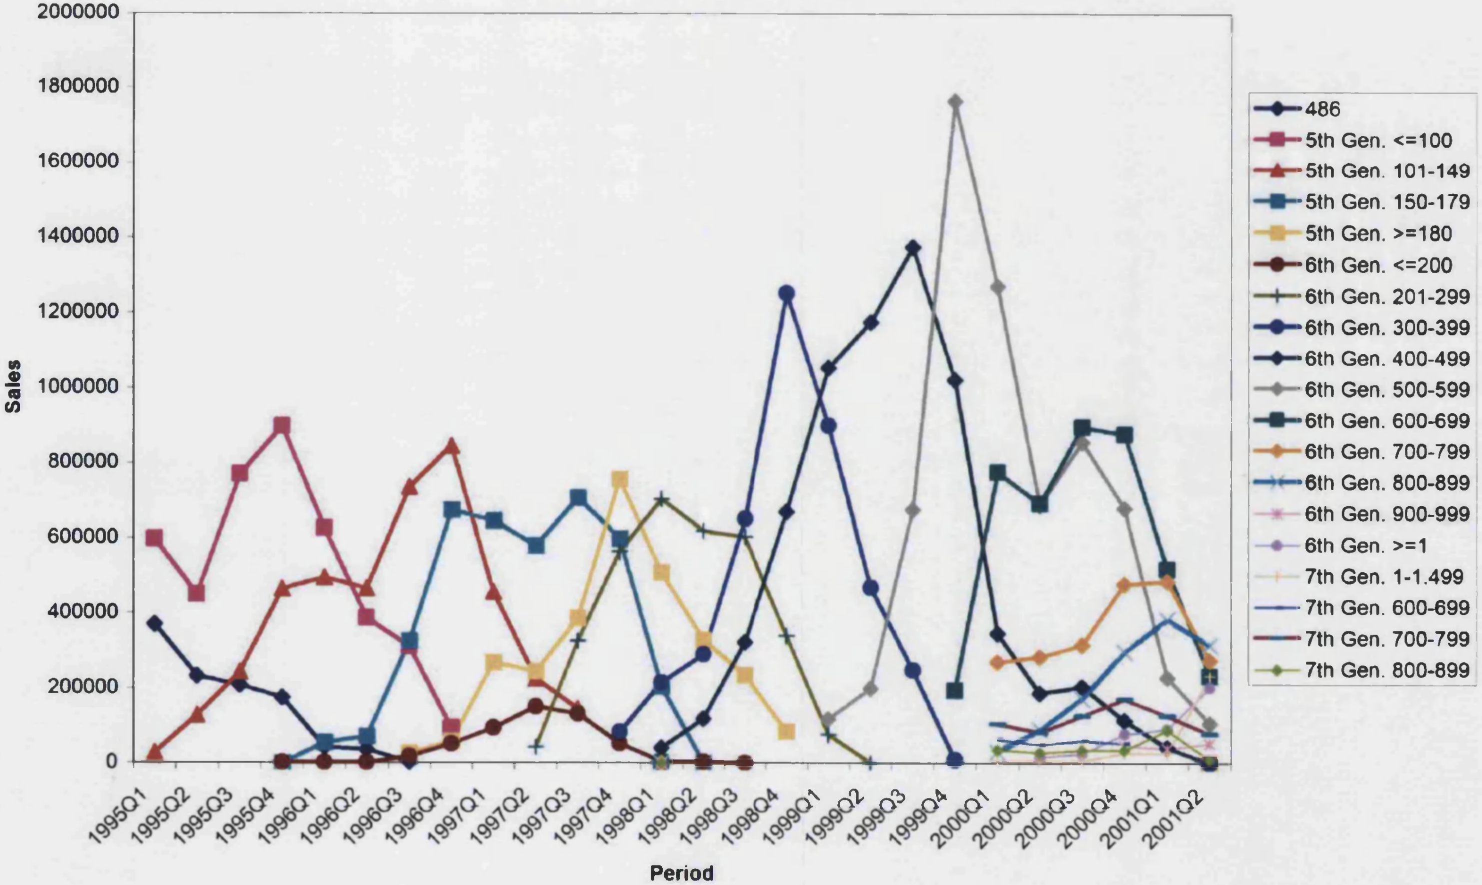

Figure 1-3: Evolution of prices by processor generation

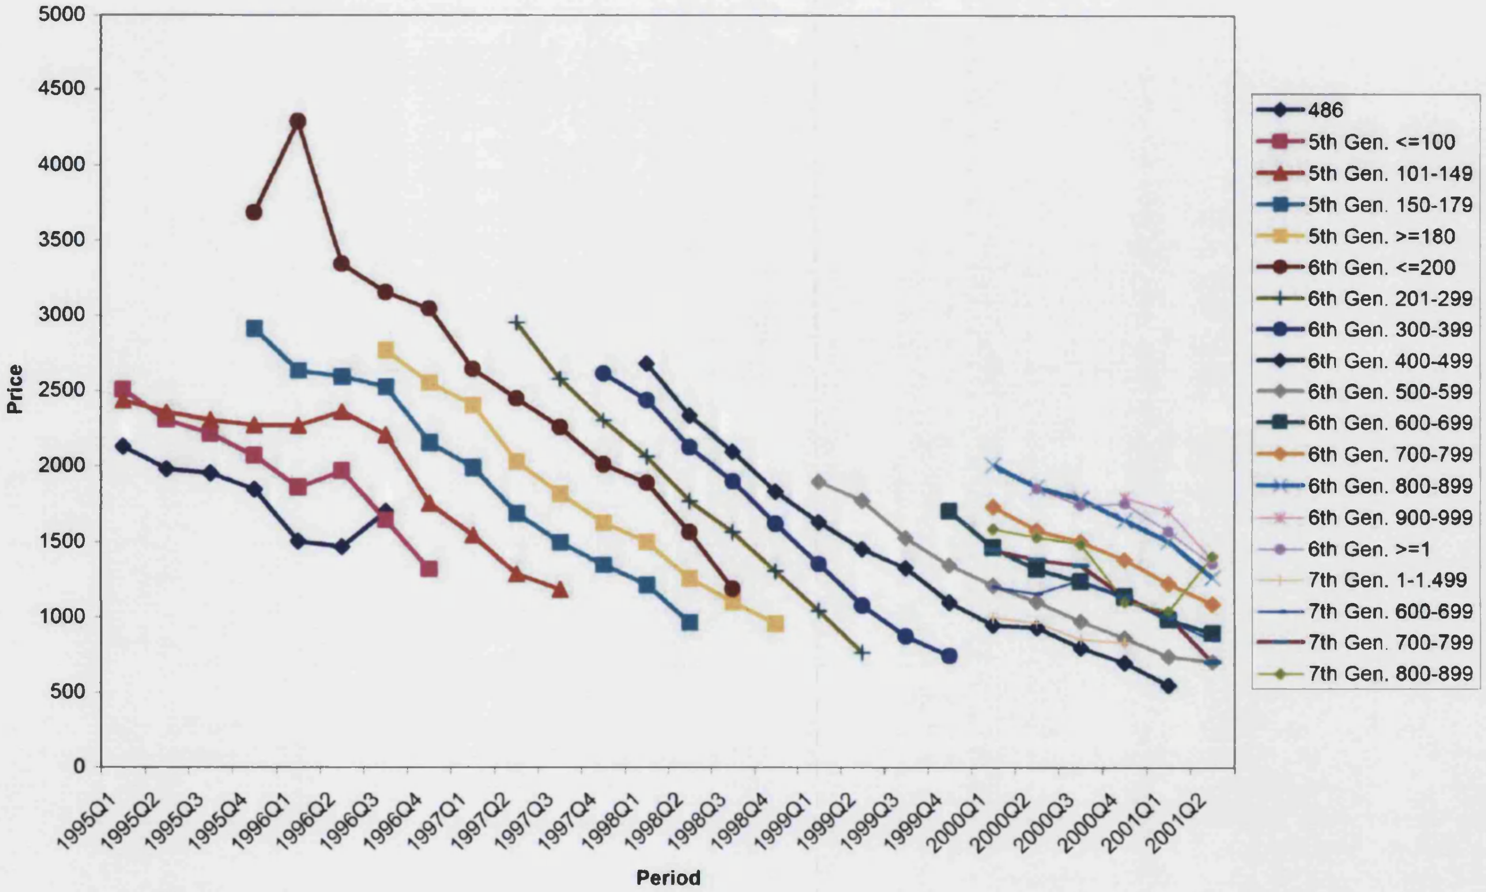

Figure 1-4: The quality window in the market for personal computers

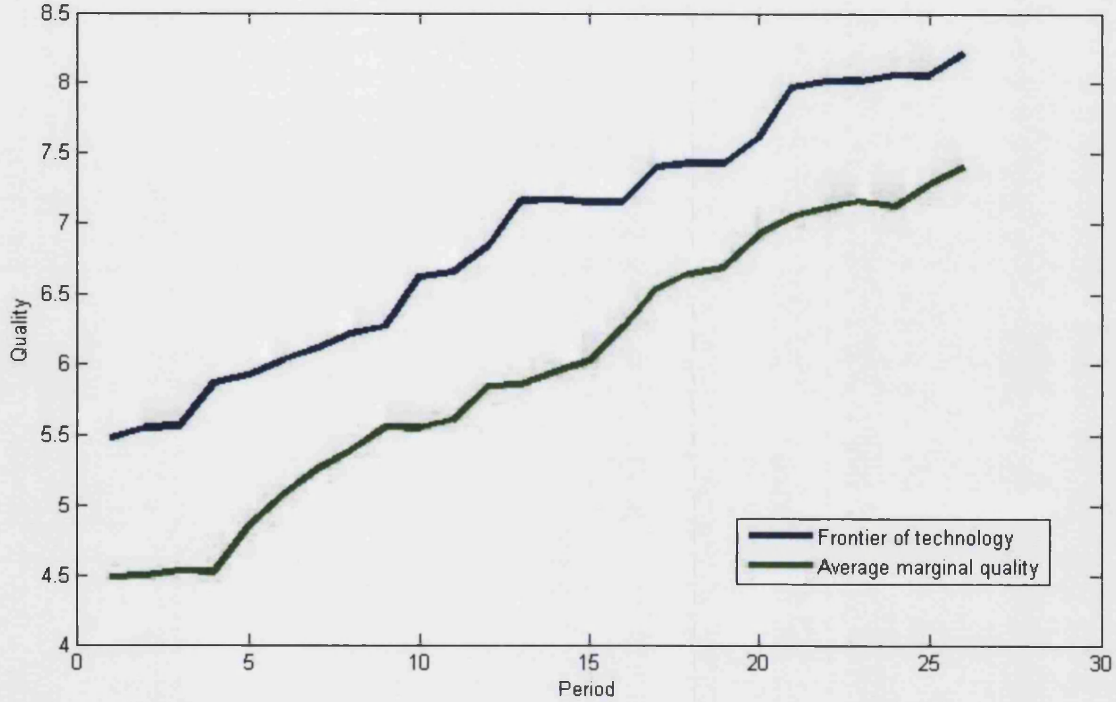

Figure 1-5: Comparison between actual and predicted quality window

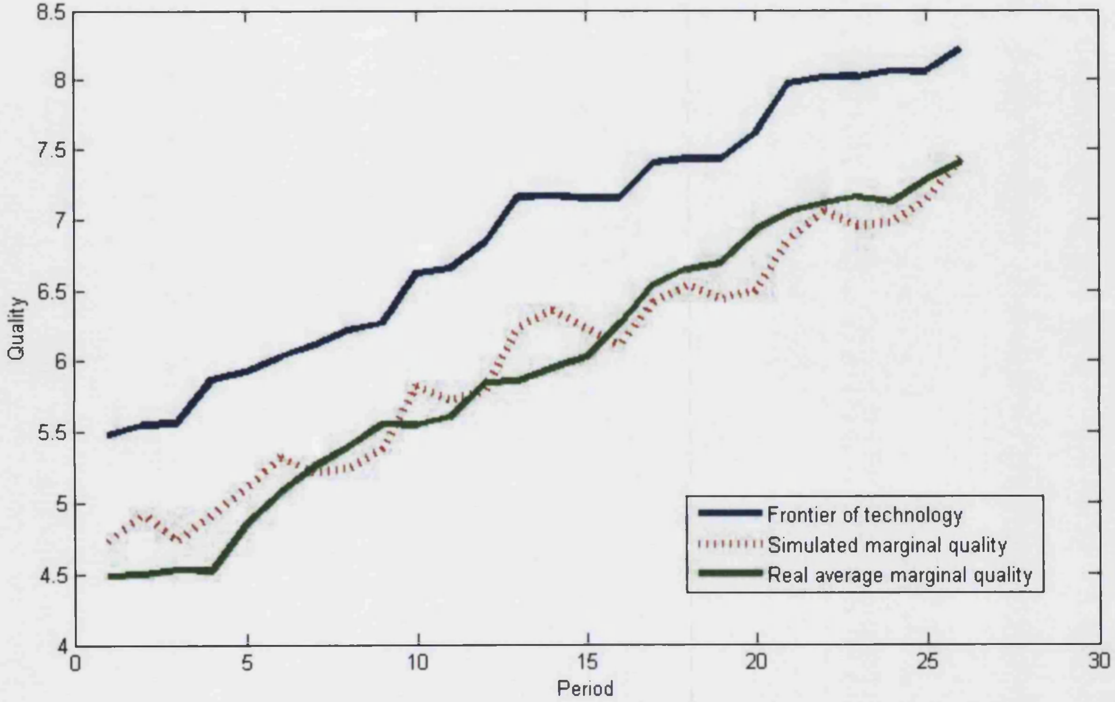

Figure 1-6: Simulation of the impact of higher willingness to pay for quality

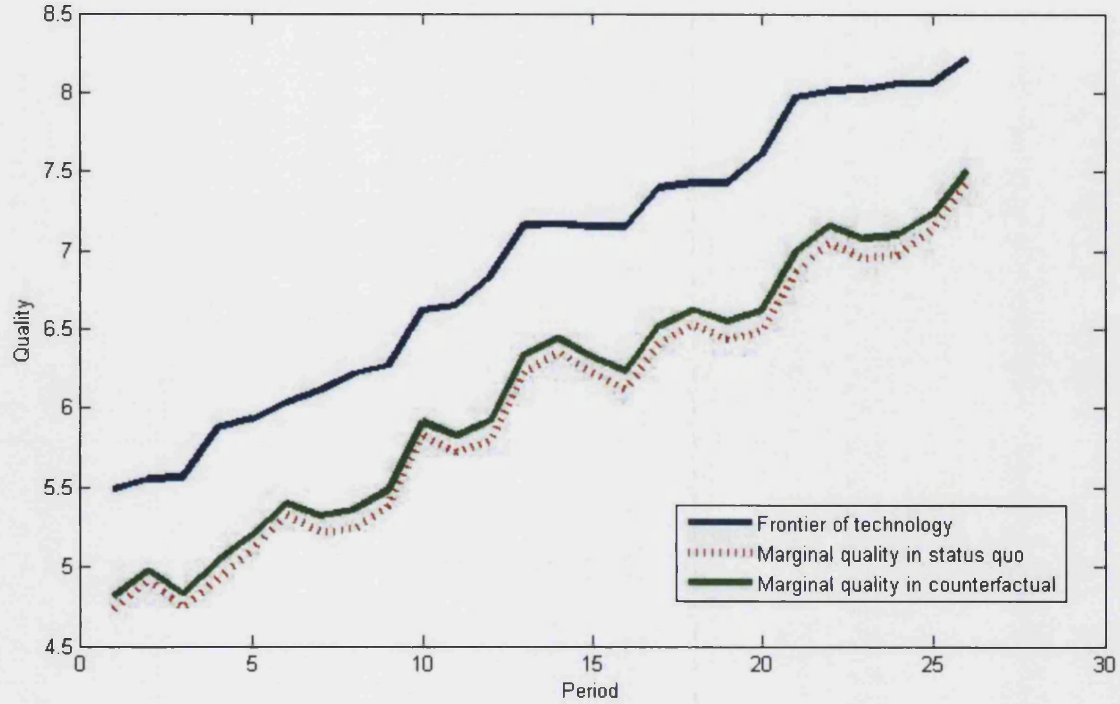

Figure 1-7: Simulation of the impact of lower cost of product variety

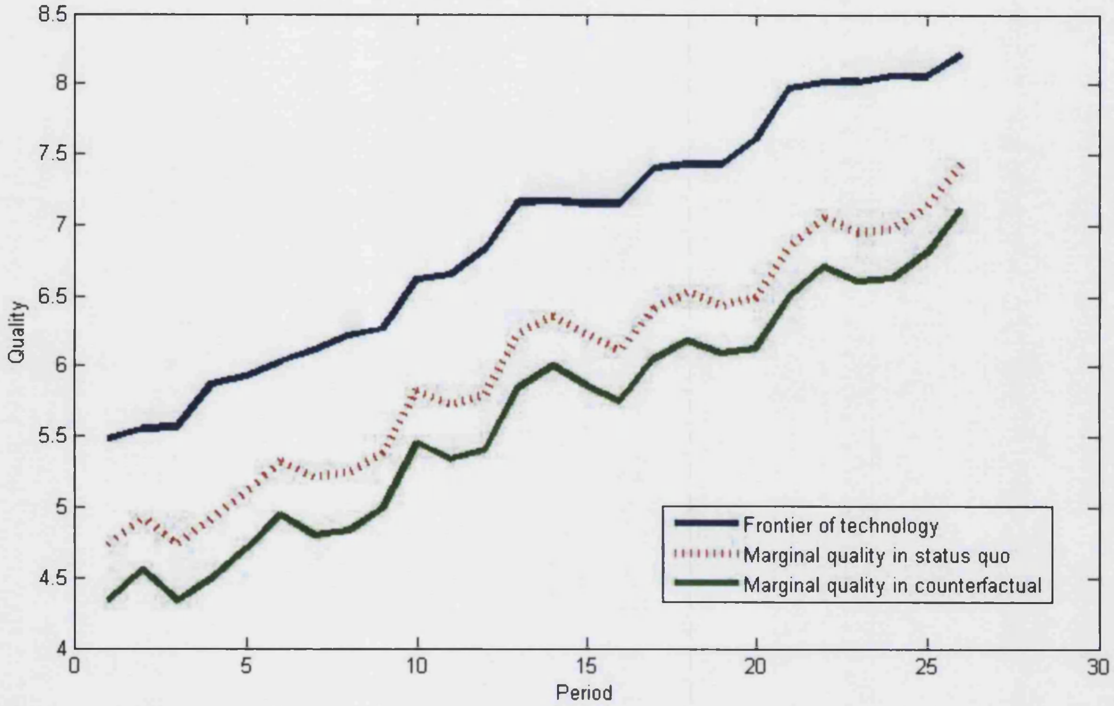

1.10 Appendix to Chapter 1

1.10.1 Derivation of the profit function and of market shares with a continuum of feasible product-qualities

Consider the (gross) profit function of firm f for a given (finite) N_j:

$$\pi_f = M \frac{\sum\limits_{q=\underline{q}_f}^{\overline{q}_f} \left[\exp\left(\alpha q + \gamma p_f\left(q\right)\right)\left(p_f\left(q\right) - \phi q\right)\right]}{\sum\limits_{f=1}^{n_f} \sum\limits_{q=\underline{q}_f}^{\overline{q}_f} \exp\left(\alpha q + \gamma p_f\left(q\right)\right) + \exp\left(\delta_0\right)}$$

Multiply both the numerator and the denominator of this expression by $\Delta q = 1/N_j$ and consider the limit for $\Delta q \to 0$:

$$\pi_f = \lim_{\Delta q \to 0} M \frac{\sum\limits_{q=\underline{q}_f}^{\overline{q}_f} \left[\exp\left(\alpha q + \gamma p_f\left(q\right)\right)\left(p_f\left(q\right) - \phi q\right)\right] \Delta q}{\sum\limits_{f=1}^{n_f} \sum\limits_{q=\overline{q}_f}^{\overline{q}_f} \exp\left(\alpha q + \gamma p_f\left(q\right)\right) \Delta q + \exp\left(\delta_0\right) \Delta q}$$

This expression can be written as

$$\pi_f = M \frac{\int_{\underline{q}_f}^{\overline{q}_f} \exp\left(\alpha q + \gamma p_f\left(q\right)\right)\left(p_f\left(q\right) - \phi q\right) dq}{\sum\limits_{f=1}^{n_f} \int_{\underline{q}_f}^{\overline{q}_f} \exp\left(\alpha q + \gamma p_f\left(q\right)\right) dq} \tag{1.10}$$

Following the same approach, it is possible to derive the market share for a (sub)set of qualities $[a,b] \subset \left[\underline{q}_f, \overline{q}_f\right]$ sold by firm f as

$$s^f_{a,b} = M \frac{\int\limits_a^b \exp\left(\alpha q + \gamma p_f(q)\right) dq}{\sum\limits_{f=1}^{n_f} \int\limits_{\underline{q}_f}^{\overline{q}_f} \exp\left(\alpha q + \gamma p_f(q)\right) dq} \tag{1.11}$$

1.10.2 Proof of proposition 1.1

Proof. (i) Each firm sets the price of each product sold. Differentiate expression 1.4 with respect to $p(q_j)$, the price of quality $q_j \in \left[\underline{q}_f, \overline{q}_f\right]$ and, according to assumption 1.2, assume that $\frac{\partial S}{\partial p(q_j)} = 0$. The first-order-condition for optimality is

$$M \frac{\exp\left(\alpha q_j + \gamma p(q_j)\right)}{S} \left[1 + \gamma\left(p(q_j) - \phi q_j\right)\right] dq = 0$$

where

$$S = \sum_{f=1}^{n_f} \int\limits_{\underline{q}_f}^{\overline{q}_f} \exp\left(\alpha q - 1 + \gamma \phi q\right) dq \tag{1.12}$$

This implies that the optimal price is

$$p(q_j)^* = \phi q_j - \frac{1}{\gamma}$$

(ii) Notice that, given optimal pricing behaviour, firm f's profit function is

$$\pi_f = \frac{M}{S} \int\limits_{\underline{q}_f}^{\overline{q}_f} \exp\left(\alpha q - 1 + \gamma \phi q\right) \left(-\frac{1}{\gamma}\right) dq - F\left(\overline{q}_f - \underline{q}_f\right) \tag{1.13}$$

Equilibrium product lines must satisfy two conditions: each firm must sell a product line that maximises its profits given the aggregate index of market output S, and the aggregate index of market output must be consistent with the choices made by all firms. Assuming that $\frac{\partial S}{\partial \underline{q}_f} = 0$, the first-order-condition that pins down the optimal marginal quality $\underline{q}_f$ is

$$\frac{\partial \pi_f}{\partial \underline{q}_f} = M \exp\left(\alpha \underline{q}_f - 1 + \gamma\phi \underline{q}_f\right) \frac{1}{\gamma} + F = 0$$

which implies that

$$\underline{q}_f^* \left(S^*\right) = \underline{q}^* \left(S^*\right) = \frac{1 + \ln\left(-\frac{\gamma F S^*}{M}\right)}{\alpha + \gamma\phi} \tag{1.14}$$

The second-order condition for a maximum is that

$$\frac{\partial^2 \pi_f}{\partial \underline{q}_f^2} = M \exp\left(\alpha \underline{q}_f - 1 + \gamma\phi \underline{q}_f\right) \frac{1}{\gamma} (\alpha + \gamma\phi) < 0$$

which is verified under assumption 1.1: $(\alpha + \gamma\phi) > 0$ because $\gamma < 0$.

Consider now the derivative of the profit function with respect to the highest quality sold $\overline{q}_f$:

$$\frac{\partial \pi_f}{\partial \overline{q}_f} = -M \exp\left(\alpha \overline{q}_f - 1 + \gamma\phi \overline{q}_f\right) \frac{1}{\gamma} - F$$

This expression is zero for $\overline{q}_f = \underline{q}^*$ and, if $(\alpha + \gamma\phi) > 0$, it is negative for any $q < \underline{q}^*$ and positive for any $q > \underline{q}^*$. Therefore there is a corner solution for a maximum: $\overline{q}_f^* = \overline{Q}$.

In equilibrium, it must be the case that the index of market output S^* is consistent with firms' decisions, i.e. that

$$S^* = n_f \int_{\underline{q}^*}^{\overline{Q}} \exp\left(\alpha q - 1 + \gamma \phi q\right) dq \tag{1.15}$$

which follows from the definition of S^*. Expression 1.6 derives from the combination of 1.14 and 1.15. ∎

1.10.3 Derivation of the estimating equation of demand

The market share of product j (defined as the sales of product j divided by total sales) is given by the standard expression:

$$sh_j = \frac{\exp\left(\alpha q_j + \gamma p_j\right)}{\sum_{i=1}^{n} \exp\left(\alpha q_i + \gamma p_i\right)}$$

The natural logarithm of product j's market share is equal to:

$$\ln\left(sh_j\right) = \left(\alpha q_j + \gamma p_j\right) - \ln\left(\sum_{i=1}^{n} \exp\left(\alpha q_i + \gamma p_i\right)\right)$$

The average of the natural logarithm of all products' market shares is:

$$\begin{aligned} \overline{\ln(sh)} &= \frac{1}{n_j}\sum_{i=1}^{n} \ln\left(sh_i\right) = \frac{1}{n}\sum_{i=1}^{n}\left(\alpha q_i + \gamma p_i\right) - \ln\left(\sum_{i=1}^{n} \exp\left(\alpha q_i + \gamma p_i\right)\right) = \\ &= \alpha\overline{q} + \gamma\overline{p} - \ln\left(\sum_{i=1}^{n} \exp\left(\alpha q_i + \gamma p_i\right)\right) \end{aligned}$$

where$n$ is the number of products, $\overline{q} \equiv \frac{1}{n_i}\sum_{i=1}^{n} q_i$ is the average quality of the products in the market and $\overline{p} \equiv \frac{1}{n_i}\sum_{i=1}^{n_j} p_i$ is their average price.

Consider now the expression $\left[\ln(sh_j) - \overline{\ln(sh)}\right]$:

$$\begin{aligned}\ln(sh_j) - \overline{\ln(sh)} &= (\alpha q_j + \gamma p_j) - (\alpha\overline{q} + \gamma\overline{p}) = \\ &= \alpha(q_j - \overline{q}) + \gamma(p_j - \overline{p})\end{aligned}$$

Given market shares calculated on the basis of total sales and quality and price information this expression can be used to estimate the parameters α and γ.

1.10.4 The impact of a change in the marginal utility of quality on profits

In this section I consider the impact of a change in the marginal utility of quality on a product's profits. Profit of product-quality q_j given by expression 1.3

$$\pi(q_j) = M\frac{\exp(\alpha q_j + \gamma p_f(q_j))(p_f(q_j) - \phi q_j)}{\sum\limits_{f=1}^{n_f}\sum\limits_{q=\underline{q}_f}^{\overline{q}_f}\exp(\alpha q + \gamma p_f(q))}$$

Substitute the optimal price $p_f(q_j) = \phi q_j + \frac{1}{\gamma}$ derived in proposition 1.1 to obtain

$$\pi(q) = M\frac{\exp(\alpha q_j - 1 + \gamma\phi q_j)\left(-\frac{1}{\gamma}\right)}{\sum\limits_{f=1}^{n_f}\sum\limits_{q=\underline{q}_f}^{\overline{q}_f}\exp(\alpha q - 1 + \gamma\phi q)}$$

Consider the derivative of $\pi_f(q_j)$ with respect to α:

$$\frac{\partial\pi_f(q_j)}{\partial\alpha} = A\left[q_j\sum_{f=1}^{n_f}\sum_{q=\underline{q}_f}^{\overline{q}_f}\exp(\alpha q - 1 + \gamma\phi q) - \sum_{f=1}^{n_f}\sum_{q=\underline{q}_f}^{\overline{q}_f} q\exp(\alpha q - 1 + \gamma\phi q)\right]$$

where $A = \frac{M\exp(\alpha q_j - 1 + \gamma\phi q_j)}{(\alpha+\gamma\phi)\sum_{f=1}^{n_f}\sum_{q=\underline{q}_f}^{\overline{q}_f}\exp(\alpha q - 1 + \gamma\phi q)} > 0$ under assumption 1.1 whereby $(\alpha + \gamma\phi) > 0$.

Therefore, $\frac{\partial \pi_f(q_j)}{\partial \alpha} > 0$ if and only if

$$q_j > \frac{\sum_{f=1}^{n_f}\sum_{q=\underline{q}_f}^{\overline{q}_f} q\exp(\alpha q - 1 + \gamma\phi q)}{\sum_{f=1}^{n_f}\sum_{q=\underline{q}_f}^{\overline{q}_f} \exp(\alpha q - 1 + \gamma\phi q)}$$

The threshold is meaningful since, as long as $\min\left\{\underline{q}_f\right\} < \max\left\{\overline{q}_f\right\}$,

$$\min\left\{\underline{q}_f\right\} < \frac{\sum_{f=1}^{n_f}\sum_{q=\underline{q}_f}^{\overline{q}_f} q\exp(\alpha q - 1 + \gamma\phi q)}{\sum_{f=1}^{n_f}\sum_{q=\underline{q}_f}^{\overline{q}_f} \exp(\alpha q - 1 + \gamma\phi q)} < \max\left\{\overline{q}_f\right\}$$

since

$$\frac{\sum_{f=1}^{n_f}\sum_{q=\underline{q}_f}^{\overline{q}_f} q\exp(\alpha q - 1 + \gamma\phi q)}{\sum_{f=1}^{n_f}\sum_{q=\underline{q}_f}^{\overline{q}_f} \exp(\alpha q - 1 + \gamma\phi q)} < \frac{\sum_{f=1}^{n_f}\sum_{q=\underline{q}_f}^{\overline{q}_f} \max\left\{\overline{q}_f\right\}\exp(\alpha q - 1 + \gamma\phi q)}{\sum_{f=1}^{n_f}\sum_{q=\underline{q}_f}^{\overline{q}_f} \exp(\alpha q - 1 + \gamma\phi q)} = \max\left\{\overline{q}_f\right\}$$

and

$$\frac{\sum_{f=1}^{n_f}\sum_{q=\underline{q}_f}^{\overline{q}_f} q\exp\left(\alpha q-1+\gamma\phi q\right)}{\sum_{f=1}^{n_f}\sum_{q=\underline{q}_f}^{\overline{q}_f} \exp\left(\alpha q-1+\gamma\phi q\right)} > \frac{\sum_{f=1}^{n_f}\sum_{q=\underline{q}_f}^{\overline{q}_f} \min\left\{\underline{q}_f\right\}\exp\left(\alpha q-1+\gamma\phi q\right)}{\sum_{f=1}^{n_f}\sum_{q=\underline{q}_f}^{\overline{q}_f} \exp\left(\alpha q-1+\gamma\phi q\right)} = \min\left\{\underline{q}_f\right\}$$

Chapter 2

A dynamic model of optimal product lines in the presence of adjustment costs

2.1 Introduction

In Chapter 1, I developed a model of endogenous product lines to explain the range of product-qualities sold in the market for personal computers as the outcome of firms' optimal decisions, and I studied the relationship between consumers' willingness to pay for quality and equilibrium product variety. The analysis relied on a static model which posited that firms could costlessly modify their product lines and therefore that they could immediately adapt the range of product-qualities sold to evolving technological and market conditions.

In general, however, adding new products to and withdrawing products from a firm's product line is not costless. For instance, introducing new products often entails significant R&D outlays and additional sunk costs to bring the developed product to the market. Similarly, withdrawing products from a firm's product line may also entail sunk "adjustment" costs, e.g. to rearrange the purchase of inputs, the production line, distribution arrangements and to update promotional material. Yet, despite the vast body of research on pricing and investment decisions in the presence of adjustment costs, and despite the importance of firms' product-line choices, I am not aware of any empirical study of the costs that firms must bear to modify their product lines.

In this Chapter, I develop a structural model of product-line decisions in the presence of sunk product-line adjustment costs which I apply to the market for personal computers. I use the model to derive an empirical estimate of the adjustment costs of adding products to and withdrawing products from a firm's product line and to study the impact that adjustment costs have on firms' profitability.

From the perspective of an individual firm, the existence of (fixed) sunk adjustment costs has both a direct and an indirect impact on profitability. The direct impact is the outlay that the firm incurs every time it modifies its product line. The indirect effect on profitability derives instead from the impact that adjustment costs have on the product line optimally sold by the firm and on the product lines sold by its competitors. When changes to the product line entail a significant fixed sunk cost, the firm does not optimally modify the products it sells every period to adapt its portfolio to changing technological and market conditions, but only when the (expected) benefits of doing so exceed the adjustment cost. This product-line rigidity implies that, in general, the set of products sold by the firm each period is sub-optimal in the sense that it does not maximise period profits (even though it is optimal in a dynamic perspective). In equilibrium, however, it is not clear, *a priori*, what the impact of adjustment costs on a firm's profitability is, because adjustment costs affect not only a firm's product line but also the product lines (optimally) sold by rival firms.

The model that I present in this Chapter is very similar to the one that I discussed in Chapter 1, and it retains much of its structure. The main difference is that, while in the static model of Chapter 1 I assumed that product-line adjustments were costless, I now consider the case in which changes to a firm's product-line entail a fixed sunk cost.

From a modelling perspective, when firms incur a cost to modify their product lines, product-line decisions become dynamic since firms need to consider when it is optimal to modify their product portfolios. Dynamic considerations enrich the analysis of product-line decisions but are also the source of considerable analytical complexity, which has to be dealt with when characterising the solution of the model and when estimating its structure.

First of all, because of the complexity of the economic environment that I consider, optimal product-line decisions cannot be described analytically as in the model that I presented in Chapter 1. In line with the literature on dynamic games of industry evolution, I therefore rely on numerical methods to characterise the equilibrium of the model. In order to reduce the dimensionality problem that stems from dealing with endogenous market structure in a differentiated market with many competitors, I rely on a model in which the impact of market structure on any firm's profits is completely summarised by an aggregate index of market output. Each firm considers only this aggregate index of market output when choosing its actions rather than the detailed description of the entire set of products sold by each firm. This approach, which is similar in spirit to the one that Melnikov (2000) proposed to estimate (logit) dynamic models of demand for durable goods, reduces the dimensionality of the state space substantially and is the source of significant computational savings.

In addition, inference of some of the structural parameters of the model is based on restrictions that derive from the model's dynamic equilibrium. An obvious nested estimation algorithm would require solving for the equilibrium of the model for each possible parameter value, which would be computationally demanding. Instead, I estimate these parameters following the two-step estimation approach developed by Bajari, Benkard and Levin (2007). The idea behind this estimator is that, if it is possible to recover an estimate of the policy functions from the data, the relevant parameters can be estimated through a simulation procedure as those that make the observed policy functions optimal. This approach greatly simplifies the empirical estimation of complex dynamic models by avoiding the need to compute the equilibrium of the model even once.

After estimating the structure of the model, I study the impact that adjustment costs have on profitability by simulating the evolution of firms' product lines over a period of 6 years (to match the period covered by my dataset) both in the presence of estimated adjustment costs and in the hypothetical counterfactual scenario in which there are no adjustment costs. The comparison of the present value of firms' simulated profits in the two scenarios suggests that, in equilibrium, the existence of adjustment costs decreases firms' profits.

This Chapter is organised as follows. Section 2.2 provides an overview of the literature to which this study is related. Section 2.3 describes some of the features of product-line adjustments in the market for personal computers that motivate the analysis. Section 2.4 presents the model. Section 2.5 describes the estimation strategy and results. The simulation of the model and the counterfactual analysis is discussed in section 2.6 and final conclusions follow.

2.2 Related literature

The economic literature on adjustment costs is large and it spans a variety of research fields. For instance, adjustment costs have been studied in the context of investment decisions (Abel and Eberly (1994), Caballero and Engel (1999)), labour demand (Caballero and Engel (1993)), pricing decisions (Caplin and Spulber (1987)) and the purchase of durable-goods (Lam (1991), Eberly (1994), Attanasio (2000)). However, to the best of my knowledge, this is the first empirical study that considers product-line adjustment costs.

From a modelling perspective, this study is related to the literature on structural dynamic models of industry evolution. Ericson and Pakes (1995) and Pakes and McGuire (1994) laid the foundations for the study of dynamic oligopolistic games of industry evolution. The development of these seminal contributions into a rich body of literature, however, has been limited by the significant computational difficulties that arise even when dealing with simple economic environments.

Gowrisankaran and Town (1997) study a dynamic game in the hospital industry, which they estimate with a computationally intensive nested GMM algorithm and use to analyse the impact of government policies. Benkard (2004) develops a multi-agent dynamic model of the commercial aircraft industry and studies industry pricing, industry performance, and optimal industry policy in presence of learning-by-doing. Besanko and Doraszelski (2004) study capacity dynamics. Dunne et al. (2006) apply the dynamic model of entry and exit developed by Pakes, Ostrovsky and Berry (2004) to the study of entry and exit of dentists and chiropractors in geographic markets. Ryan (2007) studies the welfare costs of environmental regulation on the US Portland cement industry. He recovers the entire cost structure of the industry, including the distribution of sunk entry costs and adjustment costs of investment using the same two-step estimator that I employ in this study.

Unlike most of these papers, I do not consider an oligopolistic dynamic game but a dynamic model of endogenous product-quality choice in a monopolistically competitive market. In addition, I address the problem of the dimensionality of the state space by relying on a model in which the impact of market structure of a firm's profits can be entirely summarised by a simple scalar, the "aggregate index of market output". Firms do not need to consider the detailed description of all products sold in the market but only the (expected evolution of the) aggregate index of market output when choosing optimal product lines. Melnikov (2000) used a similar approach to develop and estimate a dynamic demand model where consumers' optimal intertemporal choices can be fully described on the basis of the evolution of an index that summarises all the products in the market.

Recently, a number of contributions have significantly developed the tools for the econometric estimation of dynamic games, building upon the foundations set by the literature on the estimation of single-agent problems (Hotz and Miller (1993), Hotz et al. (1993), Rust (1994)). In this paper, I will discuss and I will use the estimating approach proposed by Bajari, Benkard and Levin (2007). Other approaches to the estimation of dynamic games have been recently proposed by Pakes, Ostrovsky and Berry (2007), by Aguirregabiria (2007) and by Pesendorfer and Schmidt-Dengler (2004).

The paper that is closest to this study is the one by Carranza (2006) who also considers a dynamic model of endogenous quality choice in a monopolistically competitive market, with an application to the market for digital cameras. Like Carranza (2006) I also develop a model of endogenous quality formation but I introduce product-line adjustment costs and consider endogenous product exit. Moreover, I use a radically different estimation approach to recover the structural parameters of the model, which does not require the solution of the dynamic equilibrium and is therefore significantly less computationally intensive.

By considering endogenous product variety in a dynamic perspective, this paper is also related to the empirical literature on product survival which has been developed by Stavins (1995), Greenstein and Wade (1998), Genakos (2005) and de Figueiredo and Kyle (2006). All these studies are based on reduced-form product entry and exit models. By contrast, in this study I am able to explicitly relate product survival to the fundamental structure of demand and supply and therefore I can potentially study how product exit is affected by changes to the economic environment.

2.3 Product-line adjustments in the market for personal computers

As I discussed in Chapter 1, the market for personal computers is a dynamic environment characterised by an ongoing process of creative destruction. Technological progress allows firms to introduce personal computers of better and better quality in the market. In turn, product innovation affects the viability of personal computers of lower quality, which are eventually withdrawn from the market in a relatively ordered sequence generating a pattern that can be described as a "moving quality window".

In Chapter 1, I studied the evolution of the aggregate range of qualities sold in the market assuming that firms could modify their product portfolios without incurring in any sunk cost and therefore that they could immediately adapt their product lines to changing technological and market conditions. However, the relative smooth evolution of the aggregate quality window that I considered in Chapter 1 conceals the pattern of product-line adjustments that takes place at the firm-level.

Figure 2-1 shows the evolution of the product lines of the 9 manufacturers over the period 1995Q1-2001Q3. In the graph, quality is proxied by the natural logarithm of the benchmark of the PC's processor. As I discussed in Chapter 1, the performance of the CPU is arguably the most important determinant of a PC's quality and it is strongly correlated with other characteristics such as the RAM and the size of the hard-disk and therefore it is a good proxy for a PC's overall quality. The Figure shows that PC manufacturers tend to sell broadly similar ranges of qualities but that there are differences, especially at the bottom of the product lines. This variance is described more precisely in Table 2.1 and in Table 2.2, where I report some descriptive statistics of the highest and lowest qualities sold by firms.

Figure 2-2 disentangles the evolution of the product lines of the largest manufacturers and shows more clearly the nature of each firm's changes to the quality window. Despite the significant pace of technological change in this industry, a typical firm does not introduce or withdraw products every quarter. On average, a firm withdraws obsolete qualities from the market only once every 2.5 quarters and it introduces new products only once every 2.2 quarters. In some cases, firms do not withdraw product-qualities from their product lines for more than a year. This relative product-line rigidity is suggestive of the existence of fixed sunk adjustment costs: since modifying the product portfolio is not costless, firms do not optimally withdraw (or introduce) products every quarter, but only when the expected benefits of doing so exceed the costs.

However, there are potentially other possible explanations for this pattern of product-line adjustments that are not related to the existence of adjustment costs. For instance, I consider a one-dimensional measure of quality and therefore I may not be able to capture changes to a firm's product line that affect product characteristics other than the speed of the CPU. However, I have verified that a product's characteristics remain constant after its introduction and therefore that I do not underestimate the frequency of changes to the marginal qualities sold by firms.

The relative rigidity of firms' product lines may also reflect the discreteness of available CPUs speeds rather than the existence of adjustment costs. I have considered whether this discreteness is driving the pattern that I observe by comparing the product-line changes that I observe to the set of feasible quality locations. I have found that when firms withdraw obsolete products from their product lines, they do so with a relative degree of lumpiness: they do not simply adjust to the next feasible quality location. This suggests that the behaviour that I observe does not simply derives from the potential rigidity imposed by a discrete choice set.

Finally, the evolution of PC manufacturers' product line may simply reflect the choices of important upstream manufacturers such as Intel: obsolete products may be withdrawn from the market simply because relative obsolete inputs become unavailable. The different inventory held by different manufacturers may drive the variation observed across different firms. However, even though the manufacturers are undoubtely subject to common technological shocks, the differences in firms' product lines and in firms' changes to product lines described by Figure 2-1 do not appear to be consistent with the idea that the evolution of firms' product lines follows passively exogenous changes that take place in upstream input markets. I have not been able to find discontinuance information of Intel products but it should be noted that, in the period considered, Intel did not differentiate between desktop and laptop processors and, even though a PC processor became obsolete for desktops, it normally continued to be supplied for use in laptop computers and (even for a longer time) for embedded applications. For instance, in August 2007, Intel Pentium MMX processors, despite being completely obsolete for personal computers, are still produced by Intel and used in embedded applications. In the model that I present in this Chapter, product exit decisions are the result of non-trivial economic choices taken by PC manufacturers and do not simply reflect the availability of relatively obsolete input components.

2.4 The model

The model describes a discrete-time industry with endogenous product entry, exit, and price-setting, where firms choose strategies in order to maximize the expected discounted value of their net future profits given their information set.

Each period $t = 1, 2, ..., \infty$, there are N consumers. At time t, each consumer can purchase one of the products offered in the market or an outside option. The utility of consumer i from purchasing a product of quality q_j sold by firm f at price p_{jf} is specified as

$$u_{ijf} = b_f + \alpha q_j + \gamma p_{jf} + \varepsilon_{ijf} = \delta_{jf} + \varepsilon_{ijf}$$

where b_f is a characteristic common to all product sold by firm f and ε_{ijf} is an i.i.d. extreme-value distributed error term. It is reasonable to assume that $\alpha > 0$ and $\gamma < 0$. Utility can be decomposed into the mean-utility term $\delta_{jf} \equiv b_f + \alpha q_j + \gamma p_{jf}$ and the stochastic component ε_{ijf}. Let δ_0 be the mean utility derived from the consumption of the outside option.

In the context of this dynamic model, the advantages of this simple logit demand system are substantial because the approach that I will use to reduce the dimensionality of the state space relies on the property of the logit demand system whereby demand for a product (and therefore its profits) depends only on the characteristics of the product and on the aggregate index of market output. Summarising market structure by means of a simple aggregate index of market output would have not been possible with more complex demand systems such as the random-coefficient logit (Berry, Levinsohn and Pakes, 1995), the (localised) pure-characteristic demand model (Berry and Pakes, 2007) or the hedonic demand model developed by Bajari and Benkard (2005).

There are n firms, indexed by $f = 1, ..., n$. Every period t, each firm sells a product line with no holes, which includes all feasible product-qualities in the interval $\left[\underline{q}_{ft}, \overline{q}_{ft}\right]$. As in Chapter 1, I assume that there is a large number of feasible qualities, so that in the limit the quality space can be considered to be continuous for the purpose of studying firms' product line decisions. A firm's product line can therefore be fully characterised by the highest quality and by the lowest quality it sells, i.e. by $\overline{q}_{ft}$ and by the "marginal quality" $\underline{q}_{ft}$.

The product-qualities sold by any firm are constrained by exogenous technological conditions. In particular, I assume that firms can sell any quality between 0 and the exogenous frontier of technology $\overline{Q}_t$. I assume that the frontier of technology evolves according to the function $\overline{Q}_{t+1}\left(\overline{Q}_t, \phi_t; \lambda_{\overline{Q}}\right)$ which is parametrised by $\lambda_{\overline{Q}}$.

All the products sold by a firm at time t share the characteristic b_{ft}, which is a vendor-specific shock that is observed each period only by firm f in advance of product-line and pricing decisions and is therefore private information. I assume that the brand vendor shock is a random variable with cumulative distribution B, and has mean equal to 0 and variance equal to σ_b^2.

A product of quality q_j is produced at constant marginal cost

$$c_t(q_j) = \phi_t q_j$$

This specification of technology implies that the marginal cost of production is linearly increasing in the quality of the product, i.e. that there are constant returns to quality. I allow the marginal cost of quality ϕ to change over time and I assume that ϕ_{t+1} evolves deterministically according to the function $\phi_{t+1}\left(\phi_t, \overline{Q}_t; \boldsymbol{\lambda}_\phi\right)$ which is parametrised by $\boldsymbol{\lambda}_\phi$.

Production also entails the "fixed" cost $F\left(\underline{q}_f, \overline{q}_f\right)$ that does not depend on the firm's output but that depends on the range of product-qualities it supplies. I assume that a firm that offers a wider product line incurs a higher fixed cost of production:

$$F\left(\underline{q}_f, \overline{q}_f\right) = F\left(\overline{q}_f - \underline{q}_f\right)$$

Each period t, the sequence of events is as follows:

1. firms observe the aggregate index of market activity that was realized in period $t-1$ (Ω_{t-1})
2. firms observe the available technological conditions $\left(\phi_t, \overline{Q}_t\right)$
3. each firm draws its random "brand" shock b_{ft} (which is private information)
4. each firm chooses its product line
5. each firm chooses the prices of the products it sells
6. consumers make their purchase decisions.

Given the nature of firms' product lines, the introduction and the withdrawal of new products is fully described by changes to $\overline{q}_f$ and $\underline{q}_f$. I assume that a firm incurs the fixed sunk adjustment cost $\underline{\Gamma}$ every time it modifies $\underline{q}_f$ and the sunk adjustment cost $\overline{\Gamma}$ every time it modifies $\overline{q}_f$. Given the evolution of product lines along the equilibrium path, $\underline{\Gamma}$ can be thought of as the sunk cost of withdrawing products and $\overline{\Gamma}$ as the sunk cost of adding new product-qualities to the product line.

In the particular application of this model to the market for personal computers that I consider in this Chapter, I interpret these costs as "adjustment" costs. However, the same modelling framework could also capture sunk costs of a different nature such as R&D outlays that must be incurred by a firm to develop new products of better quality. Also, even though I specify the "adjustment" cost as a fixed sunk outlay that must be incurred when products are added to or withdrawn from the product line, the model could easily be extended to account for variable sunk costs as well, a specification which could be more appropriate to model endogenous R&D investments.

Given the assumptions of the model, a firm's period expected profits (at the beginning of the period) are given by the following expression

$$\begin{aligned}
&\pi\left(\overline{q}_{ft}, \underline{q}_{ft}; \overline{q}_{ft-1}, \underline{q}_{ft-1}, b_{ft}, \phi_t, \overline{Q}_t, \Omega_{t-1}\right) = \\
&= M\frac{\int_{\underline{q}_{ft}}^{\overline{q}_{ft}} \exp\left(b_{ft} + \alpha q + \gamma p_{ft}(q)\right)\left(p_{ft}(q) - \phi_t q\right) dq}{E\left(\Omega_t | \Omega_{t-1}, \phi_t, \overline{Q}_t\right)} + \\
&-F\left(\overline{q}_{ft} - \underline{q}_{ft}\right) - \overline{\Gamma}\cdot\mathbf{1}\left(\overline{q}_{ft} \neq \overline{q}_{ft-1}\right) - \underline{\Gamma}\cdot\mathbf{1}\left(\underline{q}_{ft} \neq \underline{q}_{ft-1}\right)
\end{aligned} \tag{2.1}$$

where $\overline{q}_{ft} \leq \overline{Q}_t$, M is total sales, $\Omega_{t-1} \equiv \sum_{f=1}^{n} \int_{\underline{q}_{ft-1}}^{\overline{q}_{ft-1}} \exp\left(b_{ft-1} + \alpha q + \gamma p_{ft-1}(q)\right) dq$ is the aggregate index of market output observed by the firms at time t and $\mathbf{1}(\cdot)$ is an indicator function that takes a value equal to 1 if the condition $(\cdot)$

is satisfied and a value equal to zero if it is not.[1]

A firm's period expected (gross) profits depend on the product line it sells, $\left[\overline{q}_{ft}, \underline{q}_{ft}\right]$, on the marginal cost of quality ϕ and on the expected aggregate index of market output Ω that will be realised in the market, which fully summarises the impact of market structure on a firm's profits.[2] The frontier of technology $\overline{Q}$ constrains the changes that can be made to the firm's product line which entail a fixed sunk cost.

All the information relevant to each firm f's current and future expected payoffs available at time t is summarised by two vectors of state variables $\mathbf{s}_{ft} \equiv \left[b_{ft}, \overline{q}_{ft-1}, \underline{q}_{ft-1}\right] \in S_f$ and $\mathbf{s}_t \equiv \left[\Omega_{t-1}, \overline{Q}_t, \phi_t\right] \in S$. Vector $\mathbf{s}_{ft}$ describes firm f's product line at the beginning of the period and firm f's vendor shock which is observed (only) by firm f at the beginning of the period before product-line and pricing decisions are taken.

The state vector $\mathbf{s}_t$ describes the available technology and the structure of the market, which are common knowledge. There are two technological state variables: the frontier of technology $\overline{Q}_t$ and the marginal cost of quality ϕ_t.

The structure of the market is instead described by the aggregate index of market output Ω_{t-1}. In general, in order to characterise market structure, it is necessary to describe all products sold in the market by each firm $\Psi_t \equiv \left\{\overline{q}_{ft}, \underline{q}_{ft}\right\}_{f=1,..,n}$. However, in this model, the set of products sold in the market affects any product's profitability only through the aggregate index of market output Ω. By considering the index of market output rather than the entire set of qualities sold by each firm as the relevant market state variable, the dimension of the state space is greatly reduced: each firm does not have to keep track of the product line of each of its opponents but only of the evolution

[1] See Chapter 1 for the derivation of the gross period profit function.

[2] Notice that the aggregate index of market output depends on the realisation of firms' brand shocks as well as product-line and pricing decisions.

of the aggregate index of market output.

However, in the context of this dynamic model, using the aggregate index of market output Ω to summarise market structure Ψ is not without loss of generality. One important assumption that is necessary to underpin this simplifying approach is that, if two industry structures have the same aggregate index of market output at time $t-1$, then they must result in the same distribution of industry structures at time t. Let $P(\cdot)$ be the (conditional) cumulative probability distribution function of Ω_t. Specifically, I need to assume that

$$P(\Omega_t|\Psi_{t-1}, \phi_t, \overline{Q}_t) = P(\Omega_t|\Omega_{t-1}, \phi_t, \overline{Q}_t)$$

For example, the index of market output could be high either because there are many products in the market all with high prices or because there is a single product in the market with a low price. If these scenarios result in the same index of market output, the assumption requires that they must imply the same expectation of the next-period index of market output. A similar simplifying assumption appears in Hendel and Nevo (2003) and in Gowrisankaran and Rysman (2006).

Alternatively, this approach could be justified by assuming that a firm's product line is private information, i.e. that firms can only observe the aggregate index of market output and not the individual products sold by their rivals, which could be a reasonable assumption in market with a large number of products.

Each period, each firm f takes both pricing and product-line decisions. Firm f's actions taken at time t are fully described by the vector

$$\mathbf{a}_{ft} = \left[\{p_{ft}(q)\}_{q\in\left[\underline{q}_{ft},\overline{q}_{ft}\right]}, \overline{q}_{ft+1}, \underline{q}_{ft+1}\right] \in A_f$$

with $\underline{q}_{ft+1} \leq \overline{q}_{ft+1} \leq \overline{Q}_{t+1}$. Let $\mathbf{a}_t = (\mathbf{a}_{1t}, ..., \mathbf{a}_{nt})$ be the vector of all firms' actions taken at time t.

A Markov strategy for firm f is a function σ_f that maps each state into the action space: $\sigma_f : S \times S_f \rightarrow A_f$. A strategy profile is the vector $\boldsymbol{\sigma} = (\sigma_1, ..., \sigma_n)$ of strategies for each firm. If behaviour is given by a Markov strategy profile, firm f's expected profits given $\mathbf{s}$ and $\mathbf{s}_f$ can be written recursively:

$$V_f\left(\mathrm{s}, \mathrm{s}_f | \boldsymbol{\sigma}\right) = \pi_f\left(\sigma_f; \mathrm{s}', \mathrm{s}_f\right) + \beta \int\int V_f\left(\mathrm{s}', \mathrm{s}'_f | \boldsymbol{\sigma}\right) d\widetilde{P}\left(\mathrm{s}' | \boldsymbol{\sigma}, \mathrm{s}\right) dB(s'_f)$$

where $\widetilde{P}\left(\mathrm{s}'|\boldsymbol{\sigma}, \mathrm{s}\right)$ is the firm's perceived cumulative probability distribution $P\left(\mathrm{s}'|\boldsymbol{\sigma}, \mathrm{s}\right)$ over the state vector $\mathbf{s}$ and $B(s'_f)$ is the cumulative probability distribution of the firm's brand shock in the vector s_f.

The profile $\boldsymbol{\sigma}^*$ is a Markov equilibrium if, given the expectation on the evolution of the state vector, each firm f prefers its strategy σ^*_f to any alternative Markov strategy σ'_f :

$$V_f\left(\mathrm{s}, \mathrm{s}_f | \sigma^*_f\right) \geqslant V_f\left(\mathrm{s}, \mathrm{s}_f | \sigma'_f\right)$$

and the probability distribution $P\left(\mathrm{s}'|\boldsymbol{\sigma}^*, \mathrm{s}\right)$ that governs the evolution of s given equilibrium behaviour is consistent with the perceived probability distribution $\widetilde{P}\left(\mathrm{s}'|\boldsymbol{\sigma}^*, \mathrm{s}\right)$:

$$P\left(\mathrm{s}'|\boldsymbol{\sigma}^*, \mathrm{s}\right) = \widetilde{P}\left(\mathrm{s}'|\boldsymbol{\sigma}^*, \mathrm{s}\right)$$

In practice, this requires that the probability distribution $P_\Omega\left(\boldsymbol{\Omega}'|\boldsymbol{\sigma}^*, \mathrm{s}\right)$ that governs the evolution of the aggregate index of market output Ω is consistent with the perceived probability distribution $\widetilde{P}_\Omega\left(\boldsymbol{\Omega}'|\boldsymbol{\sigma}^*, \mathrm{s}\right)$.

Along the equilibrium path, therefore, firms take optimal pricing and product-line decisions with respect to a given distribution of future states and their optimal decisions generate industry transitions with the same distribution used in their optimisation.

In general, each firm's decision has an impact on the aggregate index of market output and therefore on its evolution and firms take this impact into account when choosing optimal behaviour. However, I consider the case in which, even though the evolution of the evolution of the aggregate index of market output depends upon the decisions taken by all firms, the impact of a single firm's decisions on the aggregate index of market output is negligible. In this case,

$$\widetilde{P}_{\Omega}\left(\boldsymbol{\Omega}'|\sigma^{*},\mathrm{s}\right)=\widetilde{P}_{\Omega}\left(\boldsymbol{\Omega}'|\mathrm{s}\right)$$

Under this assumption, I can study optimal behaviour assuming that each firm takes the aggregate index of market output and its evolution over time as given when choosing the optimal pricing and product-line strategy. This is equivalent to the assumption of "monopolistic competition" that I adopted in Chapter 1 and it implies that pricing and product-line decisions are not affected by cannibalisation and by strategic considerations. Even though this is potentially a restrictive assumption, it is not unreasonable if the level of concentration at the product level is very low, as it is the case in the market for personal computers.[3]

If firms choose prices without considering the impact on the aggregate index of market output and on its evolution, pricing decisions are static and equivalent to those that I studied in Chapter 1, where I showed that optimal prices entail a constant markup above marginal costs:

$$p_{ft}\left(q\right)=p_{t}\left(q\right)=c_{t}\left(q\right)-\frac{1}{\gamma}=\phi_{t}q-\frac{1}{\gamma} \tag{2.2}$$

[3]Cannibalisation considerations in a vertically-differentiated market subject to technological change are discussed in detail in Chapter 3.

Given this optimal price-setting behaviour, firm f's Bellman equation is

$$V\left(\overline{q}_f, \underline{q}_f, b, \Omega, \phi, \overline{Q}\right) = \max_{\underline{q}'_f, \overline{q}'_f} \left[\begin{array}{c} M \frac{\int_{\underline{q}_f}^{\overline{q}_f} \exp(\alpha q - 1 + \gamma\phi q)\left(-\frac{1}{\gamma}\right) dq}{\int \Omega' d\widetilde{P}_\Omega(\Omega'|\Omega,\phi,\overline{Q})} - F\left(\overline{q}_f - \underline{q}_f\right) \\ -\overline{\Gamma} \cdot \mathbf{1}\left(\overline{q}'_f \neq \overline{q}_f\right) - \underline{\Gamma} \cdot \mathbf{1}\left(\underline{q}'_f \neq \underline{q}_f\right) \\ +\beta \int \int V_f\left(\overline{q}'_f, \underline{q}'_f, b', \Omega', \phi', \overline{Q}'\right) d\widetilde{P}_\Omega\left(\Omega'|\Omega, \phi, \overline{Q}\right) dB\left(b'\right) \end{array} \right]$$

2.5 Estimation of the structural parameters

It is useful to organise the structural parameters of the model in three groups: $\boldsymbol{\theta}^s \equiv [\alpha, \gamma, \sigma_b]$, $\boldsymbol{\lambda}^s \equiv \left[\boldsymbol{\lambda}_{\overline{Q}}, \boldsymbol{\lambda}_\phi\right]$ and $\boldsymbol{\theta}^d \equiv \left[F, \underline{\Gamma}, \overline{\Gamma}\right]$. I do not attempt to estimate the discount factor β.

Vector $\boldsymbol{\theta}^s$ includes the structural parameters of demand (α, γ) and the standard deviation σ_b. The parameters in the vector $\boldsymbol{\lambda}^s$ characterise instead the evolution of the exogenous state variables $\overline{Q}$ and ϕ. Finally, the parameter vector $\boldsymbol{\theta}^d$ includes the fixed cost of production F and the sunk costs of adjustment $\underline{\Gamma}$ and $\overline{\Gamma}$: the "dynamic" parameters that are identified on the basis of firms' product-line changes.

In theory, it would be possible to estimate all the parameters of the model in one step by using a nested fixed point approach that extends the method used by Rust (1987) for single-agent dynamic programming. The approach entails

- computing an equilibrium to the dynamic model numerically given a parameter vector;
- using the computed values to evaluate an objective function based on the sample data; and
- finding the parameter vector that maximises the objective function.

In practice, however, the computational burden that results from the need to compute equilibria for each trial parameter value limits the practical feasibility of this algorithm.

Instead, I estimate the model sequentially. First of all, I estimate $\boldsymbol{\theta}^s$ and the parameter vector $\boldsymbol{\lambda}^s$. I then derive an estimate of $\boldsymbol{\theta}^d$ following the two-step methodology developed by Bajari, Benkard and Levin (2007) by (i) estimating policy functions and (ii) finding the set of structural parameters that rationalise the observed policies as optimal decisions.

This approach is not as efficient as the nested fixed-point algorithm since it does not make full use of the structure of the model in the estimation. However, the main advantage of this sequential estimation approach is that it does not require the solution of the equilibrium of the dynamic game even once. For computational reasons, the importance of estimating the structural parameters without having to compute the equilibrium of the dynamic model is not to be underestimated. Also, as discussed in Chapter 1, a sequential estimation approach does not enforce the equilibrium conditions that derive from the supply-side of the model in the estimation of demand and therefore consistency of the estimates of the demand parameters does not depend on supply-side equilibrium assumptions and is robust to a wide set of possible assumptions.

2.5.1 Static parameters

The "static" parameters are the marginal utility of quality α, the marginal utility of income γ, the marginal cost of quality ϕ and the standard deviation of the distribution of vendor shocks σ_b. Demand is specified as a static multinomial logit model. The estimating equation that I use is similar to the one that I discussed in Chapter 1:

$$\ln(s_{jft}) - \ln(s_{0t}) = \boldsymbol{\theta\xi} + \alpha q_{jt} + \gamma p_{jft} + \epsilon_{jft}$$

where "quality" is defined as the logarithm of the CPU benchmark and $\boldsymbol{\xi}$ is a vector of time-vendor dummies $\{\xi_{ft}\}_{f,t}$ that capture both the vendor shocks that are unobserved by the econometrician but that are observed by firms in advance of setting prices and the (time-varying) value of the outside good:[4] $\xi_{ft} = b_{ft} + \delta_{0t}$. ϵ_{jft} is an i.i.d. error term.[5] In order to calculate the market shares used in the regression, I have considered the size of the market to be equal to the number of U.S. households.[6]

For the purpose of addressing problems of correlation between the price and the error term, as explained in detail in Chapter 1, I use Canadian prices as an instrument for the prices of the personal computers sold in the United States. The results of the estimation are reproduced in Table 2.3 where IV estimates are compared to simple OLS results and where I also provide the results of the estimation in the case in which all the product characteristics available in my dataset are included in the regression.[7] Other alternative specifications of the demand system were presented and discussed in Chapter 1. I have also run a Durbin-Wu-Hausman test for endogeneity: the result suggests that it is possible to reject the hypothesis that OLS estimates are consistent at the 1% level.

[4] ξ_{ft} is a variable that takes value 1 for firm f and time t and value 0 otherwise.

[5] Notice that the time-vendor dummies could potentially also control for the effect of advertising (at the vendor level) which is unobserved by the econometrician. Controlling for advertising can be important because, if advertising occurred when prices were reduced, advertising could bias the estimate of γ upwards.

[6] Notice that, given the full set of time vendor dummies included in the regression, market size does not actually affect the estimates of the parameters α and γ.

[7] Given its large size, the vector of time-vendor dummies is not reported in the Table.

In order to disentangle the vendor shocks b_{ft} from the value of the outside good $\delta_{0t,}$ I regress the estimated time-vendor dummies ξ_{ft} on time dummies:

$$\xi_{ft} = t_1 + t_2 + ... + t_{26} + u_{ft}$$

The results are reported in Table 2.4.I use the residuals of this regression (u_{ft}) as an estimate of the vendor shocks b_{ft} and their standard deviation as an estimate of σ_b, which I have found to be equal to 1.17.

As in Chapter 1, the estimated marginal utility of income $\widehat{\gamma}$ can be used to infer the marginal cost(s) of quality on the basis of the first-order-condition for optimal pricing given by expression 2.2. The estimating equation is

$$p(q_i) + \frac{1}{\widehat{\gamma}} = \phi_t q_i + u_{it}$$

The estimates of the time-varying marginal cost of quality are reported in Table 2.5. It can be noted that the average relative price-cost margin implied by the estimated marginal utility of income is about 25%. As discussed in Chapter 1, this is consistent with the average gross profit margin of the main PC manufacturers in the period considered of about 24% that I have recovered from the firms' annual accounts.

The approach that I have adopted to the estimation of the marginal cost of quality is feasible because I have assumed that each firm ignores the impact of its pricing decisions on the aggregate index of market output, and therefore on its evolution. If pricing decisions had a significant impact on the aggregate index of market output, then optimal pricing would entail dynamic considerations, i.e. prices would not simply maximise period profits.

2.5.2 Evolution of state variables

I need to estimate the evolution of two exogenous state variables: the frontier of technology $\overline{Q}$ and the marginal cost of quality ϕ. I have specified and estimated the following regressions:

$$\overline{Q}_t = \lambda_{\overline{Q}}^1 + \lambda_{\overline{Q}}^2 \overline{Q}_{t-1} + \lambda_{\overline{Q}}^3 \phi_{t-1}$$

and

$$\phi_t = \lambda_{\phi}^1 + \lambda_{\phi}^2 \phi_{t-1} + \lambda_{\phi}^3 \overline{Q}_{t-1}$$

I have measured the frontier of technology $\overline{Q}$ by the average top quality sold by firms in the market each period.[8] In the regressions, I have used the estimates of the marginal cost of quality $\widehat{\phi}$ that I had derived earlier. The estimated parameters are presented in Table 2.6 and in Table 2.7. Note that the reported standard errors are not corrected for the sample errors in the first-stage estimates.

In addition to these exogenous technological state variables, I also estimate the evolution of the aggregate index of market output. The aggregate index of market output is an endogenous variable at the industry level but, under the assumptions of the model, it is exogenous from the perspective of an individual firm and I can treat it as such when I estimate the dynamic parameters of the model (see below). I have specified the evolution of the aggregate index of market output as follows

$$\log(\Omega_t) = \lambda_{\Omega}^1 + \lambda_{\Omega}^2 \log(\Omega_{t-1}) + \lambda_{\Omega}^3 \phi_{t-1}$$

[8] I could have defined the frontier as the maximum quality sold in the market each period. However, since I will use the exogenous evolution of the frontier of technology in the simulation of the industry equilibrium, considering the average top quality sold by firms captures better the evolution of the range of products sold in the market.

Since I do not observe the true aggregate index of market output Ω_t, I construct an estimate as

$$\widetilde{\Omega}_t \equiv \sum_{f=1}^{n} \int_{\underline{q}_{ft}}^{\overline{q}_{ft}} \exp\left(\widehat{b}_{ft} + \widehat{\alpha} q + \widehat{\gamma}\widehat{\phi}_t q\right) dq$$

where $\widehat{\alpha}, \widehat{\gamma}, \widehat{b_{ft}}$ and $\widehat{\phi}_t$ are the estimated parameters that I had previously derived. The regression results are reported in Table 2.8. Note again that the reported standard errors are not corrected for the sample errors in the first-stage estimates.

2.5.3 Dynamic parameters

The estimation of the model requires also the recovery of three "dynamic" parameters: the fixed cost of production F, the sunk cost of adding new products to the product line $\overline{\Gamma}$, and the sunk cost of withdrawing products from the product line $\underline{\Gamma}$. These parameters are "dynamic" in the sense that they are identified on the basis of firms' (dynamic) product-line decisions.

Given estimates of the parameter-vectors $\boldsymbol{\theta}^s$ and $\boldsymbol{\lambda}^s$, the fixed cost of production and the sunk costs of adjustment are inferred through a two-step procedure. In the first step, I recover an estimate of the policy functions that describe firms' product-line adjustments. In the second step, I find the dynamic parameters $\boldsymbol{\theta}^d$ that rationalise the observed policies as optimal given the theoretical model and the observed evolution of state variables. I do so by forward simulating a firm's product line when it follows the estimated policy functions, given the evolution of exogenous state variables, and by calculating the firm's simulated present-value profits as estimates of value functions. I then randomly perturbate the policy functions so as to generate different paths and different present-value payoffs. At the true parameters, the payoffs generated by the observed policies should be greater than those generated by any other set of policies.

Formally, following Bajari et al. (2007), let x index the equilibrium conditions, so that each x denotes a particular $\left(\mathrm{s}, \mathrm{s}_f, \sigma_f'\right)$ combination. The strategy profile $\boldsymbol{\sigma}$ is an equilibrium if and only if for all firms f and all states s, s_f

$$V_f\left(\mathrm{s}, \mathrm{s}_f; \sigma_f; \theta^d\right) \geqslant V_f\left(\mathrm{s}, \mathrm{s}_f; \sigma_f'; \theta^d\right)$$

Consider the difference between the value function when the optimal policy σ_f is followed and when another policy σ_f' is followed:

$$g\left(x; \theta^d, \theta^s\right) = V_f\left(\mathrm{s}, \mathrm{s}_f, \sigma_f; \theta^d, \theta^s\right) - V_f\left(\mathrm{s}, \mathrm{s}_f, \sigma_f'; \theta^d, \theta^s\right)$$

where $\boldsymbol{\sigma}$ is parametrised by θ^s.

Define the function

$$Q\left(\theta^d, \theta^s\right) = \int \left(\min\left\{g\left(x; \theta^d, \theta^s\right), 0\right\}\right)^2 dH(x)$$

where H is the distribution of the set of inequalities considered. The true parameter vector satisfies

$$Q\left(\theta_0^d, \theta_0^s\right) = 0 = \min_{\theta^d} Q\left(\theta^d, \theta_0^s\right)$$

In practice, θ^d is estimated by minimising the sample analogue of $Q\left(\theta^d, \theta_0^s\right)$.

In order to do so, I consider a random set of inequalities by randomly perturbating the estimated policy functions.[9] For each inequality, I use the forward simulation procedure to construct analogues of the V_f terms and

[9] The particular method for selecting inequalities will affect efficiency but the only requirement for consistency is that H has sufficient support to yield identification.

$\widetilde{g}\left(x;\theta^d,\theta^s_n\right)$, the empirical counterpart to $g\left(x;\theta^d,\theta^s_0\right)$ which is computed by replacing V_f with the simulated estimates.

The parameter vector θ^d is estimated as

$$\widehat{\theta^d} \equiv \arg\min_{\theta^d} \widetilde{Q}_n\left(\theta^d,\widehat{\theta^s_n}\right)$$

where

$$\widetilde{Q}_n\left(\theta^d,\widehat{\theta^s_n}\right) \equiv \frac{1}{n_I}\sum_{k=1}^{n_I}\left(\min\left\{\widetilde{g}\left(X_k;\theta^d,\widehat{\theta^s_n}\right),0\right\}\right)^2$$

and n_I is the number of simulation draws.

In order to accelerate the estimation algorithm, I decompose the profit function into a linear function of its known and unknown components as suggested by Bajari et al. (2007). The unknown components enter linearly into the payoffs of the firm in the current period and all future periods. It is therefore possible to decompose the value function into the vector of parameters and the vectors of expected discounted payoffs and actions. I thus can write the value function as

$$V_f\left(\mathrm{s},\mathrm{s}_f;\sigma_f;\theta^d\right) = W_f\left(\mathrm{s},\mathrm{s}_f;\sigma_f\right)\cdot\theta^d$$

This linearity avoids having to repeat the simulation procedure for each potential parameter value. Since W_f does not depend on θ^d, I use the forward simulation procedure once to estimate W_f and then multiply by θ^d to estimate V_f for any value of θ^d.

The first step of the estimation strategy that I have outlined entails the estimation of policy functions. In general, provided that the data is rich enough, policy functions can be estimated parametrically or non-parametrically. However, the economic environment that I consider is significantly non-stationary and therefore I do not observe the behaviour of firms at any given state repeatedly. I therefore rely on a parametric specification of firms' optimal behaviour to inform the forward-simulation procedure.

In the absence of clear theoretical results on the nature of optimal product-line decisions, the specification of policy functions necessarily involves a pragmatic approach. In general, in the presence of fixed product-line adjustment costs, the description of optimal behaviour requires the specification of two elements: (i) when the adjustment occurs and (ii) the level of the adjustment. I follow a large literature on optimal behaviour in the presence of fixed sunk adjustment costs and specify policy functions as (S, s) rules. Because of the difficulties in proving the optimality of an (S, s) rule, a pragmatic approach of this kind in the presence of fixed sunk adjustment costs was followed, for instance, by Attanasio (2000) and by Ryan (2007).

When the optimal introduction (withdrawal) of products takes the form of an (S, s) rule, policy functions can be described by the specification of a band: the adjustment takes place if and only if the firm's state is below a certain threshold s and the firm adjusts to a target level S. Both s and S are in general a function of state variables. In order to characterise optimal behaviour, it is thus necessary to specify and estimate (only) the threshold s and the target level S.

I consider separately the (S, s) rule that governs the withdrawal of product from the product line (i.e. changes to the marginal quality $\underline{q}_f$) and the (S, s) rule that underlies the introduction of new products to the product line (i.e. changes to $\overline{q}_f$).[10] Under the assumption that each threshold and target level

[10]This approach is justified by the fact that I explicitely control for the vendor brand shocks and, under the asssumption of monopolistic competition, the firm does not consider the impact that one product in its product line has on the other products it sells and by the fact that the sunk cost of introducing new products in the product line is additional to the

are observable when a firm makes an adjustment, it is possible to recover consistent estimates of the policy functions using simple OLS regressions.[11] In general, it is desirable to use a flexible high-order polynomial, so as to minimise the imposition of possibly unjustified *a priori* parametric restrictions on policy functions. However, I have found that relatively parsimonious regressions describe rather well the evolution of product-lines that I observe in the data.

The threshold for the marginal quality, $\underline{q}^s$, determines *when* a firm optimally withdraws obsolete products from the product line. In order to estimate the threshold for $\underline{q}$, I consider all those cases in which a firm changes the marginal quality it sells and I regress the firm's marginal quality before the change on the lagged market state Ω, on the technology state ϕ and on the firm's brand shock b that I previously recovered. The estimated threshold is

$$\underline{q}_t^s = \underset{(0.86)}{4.875} + \underset{(0.05)}{0.179} \log\left(\Omega_{t-1}\right) - \underset{(1.58)}{4.546}\phi_t + \underset{(0.02)}{0.023} b_t$$

In order to estimate the target level for $\underline{q}$, $\underline{q}^S$, I regress the observed marginal quality sold by a firm on the lagged index of aggregate output Ω, on the technology state ϕ and on the firm's brand shock b every period in which I observe a change of the marginal quality sold by a firm. I have estimated the target marginal quality as

$$\underline{q}_t^S = \underset{(0.73)}{4.471} + \underset{(0.04)}{0.203} \log\left(\Omega_{t-1}\right) - \underset{(1.34)}{2.234}\phi_t + \underset{(0.02)}{0.043} b_t$$

Turning now the optimal introduction of new products in the product line, I estimate the threshold $\overline{q}^s$ by considering all the instances in which a firm introduces new better qualities to its product line. In particular, if I observe

sunk cost of withdrawing products from the product line.

[11] See Ryan (2005) and Attanasio (2000).

a change in period t, I regress the firm's top-quality sold in period $t-1$, $\overline{q}_t$, on $\overline{Q}_t$:

$$\overline{q}_t^s = \underset{(0.003)}{0.97}\,\overline{Q}_t$$

I assume that when firms introduce new products to their product lines, they introduce all qualities up to the available frontier of technology, i.e. the target is $\overline{q}_{t+1} = \overline{Q}_t$.[12]

Since I deal with a non-stationary environment, caution is necessary in interpreting the policy functions that I have recovered as structural relationships. In particular, it is necessary to acknowledge that the estimated policy functions may not well capture optimal behaviour in states that are significantly different to those that I observe in the sample. I will account for this limitation of the estimated policy function, i.e. for possible out-of-sample error, in the second step of the estimation approach.

The assumption that the product-line changes that I observe in the data, as described by the estimated (S, s) rules, reflect optimal behaviour underpins the empirical strategy that I use to estimate the vector of "dynamic" parameters.

In the second-step of the estimation, I use the estimated policy functions, and the estimated processes governing the evolution of exogenous state variables, to simulate forward the evolution of firms' product lines so as to calculate the present value of simulated profits as an estimate of the value function. I have considered 30,000 simulations and for each simulation 200 random paths of brand shocks.

In order to construct estimates of the value function at different states, I have focused on the range of states that I observe in my sample.

In general, it is desirable to simulate the evolution of firms' product line for

[12]This is consistent with my model in which there are no sunk variable costs of product introduction and in which period profits are increasing in the quality of the product.

a large number of periods, so that payoffs towards the end of the time-horizon have a very low discounted present value. In order to correctly calculate the value function given optimal behaviour, it is also necessary to rely on estimated policy functions that apply in all the states that are visited in the simulated evolution of the industry. However, as I noticed above, given the non-stationarity of the economic environment that I consider, I do not think that it is reasonable to interpret the policy functions that I have estimated as structural relationships that necessarily describe well optimal behaviour also in states that are significantly different to those that I observe in my sample. The length of the industry simulation should thus be chosen so as to strike a balance between the possible out-of-sample prediction error of using the estimated policy functions and the error in estimating the value function as the present value of simulated profits in a limited number of periods only.[13] I have used a time-horizon of 30 periods to simulate the evolution of the industry. As a sensitivity check, I have verified that changing the length of the simulation does not actually materially affect the estimates obtained.

I then construct alternative paths of product lines evolution by randomly perturbating the estimated target level and the threshold for $\underline{q}$ and the estimated threshold for $\overline{q}$ by adding stochastic shocks drawn from normal distributions. I compute the present value of the simulated profits associated with these alternative policies and I use the inequality estimator described in section 2.5.3 to infer the parameter vector θ^d. Intuitively, the parameter vector is estimated as that which rationalises the estimated policy functions as optimal given the evolution of exogenous state variables.

The estimates of $\mathcal{F}$, $\overline{\Gamma}$ and $\underline{\Gamma}$ that I have obtained are reported in Table 2.9.

[13] More generally, these considerations suggest that it would be reasonable to put more weight on the information provided by state configurations that I observe in the early periods of the sample.

I do not provide the standard errors for these estimates. The main source of variance is the sample error in the first-step estimates and the estimation would require the use of subsampling or of the bootstrap. I leave the calculation of the standard errors of these dynamic parameters to future research.

The estimate of F that I have obtained is \$72.3 millions, which is higher that the estimate of \$49.9 millions that I obtained in the static model of Chapter 1. By calculating standard errors, it would be possible to assess if the difference is statistically significant. The estimated fixed cost of production is about 60% of average period gross profits.

The estimate of the sunk cost of introducing new products in the product line $\overline{\Gamma}$ is \$4.1 millions and the estimate of the sunk cost of withdrawing products from the product line $\underline{\Gamma}$ is \$2.2 millions. The estimated sunk cost of introducing new products is higher than estimated the sunk cost of withdrawing products from the product line, which is reasonable.[14] Unfortunately, however, I have not been able to identify any available source of external data that I could use as a benchmark to compare the estimates of product-line adjustment costs that I have obtained.

2.6 Simulation of industry evolution and counterfactual analysis

Having estimated its structure, I use a simplified version of the model as a tool to simulate the evolution of the industry, i.e. to replicate the product-line dynamics that I observe in my sample, and to study the impact that the existence of product-line adjustment costs has on firms' profitability. In

[14]However, without estimates of standard errors, it is not possible to assesswhether or not the difference between the two estimates is statistically significant.

particular, in the simulations that I present in this section, I focus on product exit as the only endogenous dynamic decision taken by firms: in order to ease the numerical analysis, I assume that firms introduce new better products in their product lines as soon as it is feasible to do so given the exogenous evolution of the frontier of technology. Also, I do not consider the firm-specific vendor shocks as they significantly increase the dimension of the state space.

Table 2.10 summarises the parameters of the model that I have used for the simulation and Table 2.11 describes the initial conditions that I have set, which match the states that I observe/have estimated in the first period of my sample.[15] The specification of the initial conditions is important because, given the non-stationarity of the market environment that I consider, I use the model to reproduce the evolution of the industry over a specific time horizon, starting from the same initial conditions that I observe in the market, rather than focusing on the long-run equilibrium of the industry.

Given the nature of the evolution of the exogenous state variables that I consider, the market environment that I study does not have a bounded state space. For instance, the technological frontier shifts continuously forward, without approaching an upper bound. This unboundedness would imply that innovative opportunities to improve the quality of personal computers would never fade away and therefore may not be reasonable in the very long run. In addition, it also poses a problem in solving the model numerically since this requires the specification of a bounded state space. There is no good general solution to dealing with unbounded state spaces and the approach that I have adopted is to consider a truncated but "sufficiently large" state space.

I believe that this solution is satisfactory because I am not interested in

[15] Notice that, since the model is not able to capture inter-firm variation in product-line dynamics, I have considered the simplified case in which the initial state configuration is such that all firms sell the same range of product-qualities.

the properties of the industry equilibrium in the long-run but in simulating the evolution of the industry for limited number of periods so as to match the industry evolution that I observe in my sample. Even though I consider a truncated bounded state space, as long as this state space is sufficiently larger than the state space that I would expect to encounter in practice given the estimated processes governing the evolution of the state variables over the number of periods that I am interested into, I should be able to characterise with a sufficient degree of accuracy the equilibrium of the model over this limited time-horizon.

I solve the model using a nested fixed-point algorithm to implement the "rational expectation" equilibrium in which optimal decisions are consistent with the evolution of the aggregate index of market activity generated by the optimal behaviour of all firms. The inner fixed point solves for the value function given firms' beliefs over the evolution of the index of market output. The outer algorithm derives an estimate of the evolution of the index of market output by simulating the evolution of the industry given optimal behaviour as resulting from the inner fixed point. The equilibrium is found when the simulated evolution of the index of market activity is consistent with firms' beliefs.

I have simulated the evolution of the industry for 40 periods but I only report the results of the first 26 periods, so as to match the time horizon that I observe in the data. Figure 2-3 compares the estimated quality window (continuous lines) with the average range of qualities observed in the market (dotted lines) over the period 1995Q1-2001Q2.[16]

The simulated evolution of the marginal quality $\underline{q}$ matches relatively well

[16]Notice that the bottom dotted line is an average across firms in the sample and therefore smoother than the estimated lower bound.

the evolution of the range of qualities sold in the market for personal computers, even though it tends to overestimate the average marginal quality, and it captures the discontinuous nature of product-line adjustments that characterises firms' changes to product lines. However, it should also be noted that the simulation tends to underestimate the frequency of product-line adjustments compared to what I observe in my dataset: while in my dataset I observe that an average firm modifies the marginal quality it sells once every 2.5 quarters, in the simulated path of product line evolution the adjustment takes place every 3.7 quarters.

The main benefit of estimating a structural model is the ability to simulate the effect of changes to the model primitives on equilibrium outcomes in the absence of relevant natural experiments. In particular, I have used the model to study the impact that a reduction in the level of adjustment costs (to withdraw products from the product line) has on a firm's profitability.

From the perspective of an individual firm, i.e. for a given evolution of the aggregate market state, adjustment costs negatively affect profitability in two ways. First of all, the firm must incur the outlay every time it changes its product line. In addition, in the presence of adjustment costs, the firm will sell a product line that is sub-optimal in the sense that it does not maximise period profits, even though the product line is optimal in a dynamic perspective: because of (optimal) product-line rigidity, each period the firm may be selling products that are not profitable or it may not be selling profitable products. In equilibrium, however, the impact of adjustment costs on a firm's profitability depends also on the impact that it has on the product lines sold by the firms' rivals, i.e. on the aggregate market state, and it is not clear *a priori.*

I have used the model presented in this Chapter to investigate the impact of adjustment costs of firms' profitability by simulating, over the 26-period time horizon, the products sold by a representative firm in the equilibrium when there are no adjustment costs, i.e. $\underline{\Gamma} = 0$. Figure 2-4 shows the difference between the evolution of the representative firm's product line with (dotted line) and without (continuous line) adjustment costs.[17] It is evident that, when there are no adjustment costs, the withdrawal of obsolete products occurs more frequently since firms are able to adjust their product lines every quarter to adapt to technological and market conditions (the introduction of new better products in the market and the change in the marginal cost of production). The comparison is also useful to understand the nature of product-line adjustments in the presence of adjustment costs: when a firm adjusts its product line, it generally chooses a marginal quality that is higher than the marginal quality that would maximise period profits and keeps this quality on the market for longer than would be optimal if there were no adjustment costs.

I have computed the present value of firm's profit in the status quo with adjustment costs and in the hypothetical counterfactual scenario with no adjustment costs. I have found that, in equilibrium, the existence of adjustment costs decreases (the present value of) a firm's profits by approximately $32 millions over the period considered, with adjustment-cost outlays of about $11 millions (in present value) accounting for only about one third of the overall difference in profits. This result suggests that the indirect impact of adjustment costs on firm's profitability is significantly larger than the direct impact, i.e. than adjustment costs outlays.

[17]Notice that the evolution of the marginal quality sold in the market is smoother than the one obtained in Chapter 1. The main reason is that in this dynamic model, in order to reduce the dimensionality of the state space, I have not considered changes to market size, i.e. total sales, over time.

2.7 Conclusions

In this Chapter I have developed a dynamic model of product-line decisions in which adding products to and withdrawing products from a firm's product line entails a sunk cost.

Modelling dynamic behaviour in a differentiated and technologically evolving market with a relatively large number of firms is particularly challenging because of the computational burden of solving a dynamic programming problem with a large and potentially unmanageable state space. In this Chapter I have dealt with the problem of dimensionality by relying on a model in which the impact of market structure on profits can be fully described by a simple aggregate index of market output. Firms can therefore take product-line decisions on the basis of this aggregate index of market output, rather than on the basis of the detailed consideration of all products sold in the market, as the relevant market state variable.

I have estimated the structure of this model by using the approach developed by Bajari et al. (2007), which entails recovering policy functions from the data and then finding the parameters that rationalise the estimated policy functions as optimal. This approach is not computationally demanding since it does not involve solving for the equilibrium of the dynamic model even once. In particular, I have derived estimates of the sunk costs that firms incur when adding products to and withdrawing products from their product lines. To the best of my knowledge, this is the first structural empirical study to derive estimates of product-line adjustment costs.

I have used the model to simulate firms' product line decisions and the endogenous evolution of the "quality window" in the market for personal computers. The model matches relatively well the evolution of the range of qualities and the rigidity of firms' product lines in the market for personal computers.

I have also used this structural model to study the impact that product-line adjustment costs have on a firm's profitability by simulating the impact of a change in the level of adjustment costs. I have found that the existence of adjustment costs decreases firms' profits in equilibrium and that the main impact is indirect: adjustment costs outlays as such account for only one third on the reduction in (the present value of) profits.

The model can be improved and developed in a number of directions.

First of all, the model would significantly benefit from the use of a demand system that is better suited than the logit to deal with markets in which there is a large number of products. However, the potential advantages of a different demand system need also to be weighted against any potential disadvantages for the analysis of the supply-side of the model. In particular, in this model I have been able to deal with a very complex market environment because I have used a specification in which the impact of market structure on a firm's profitability can be summarised by a simple scalar. This is a property that crucially depends on the choice of the demand system.

At present, the simulation that I have presented does not match the inter-firm product-line variation found in the data. This variation could be driven by the idiosyncratic vendor shocks that I have considered in the model, but their consideration in the simulation significantly increases the state space and raises computational problems that need to be addressed.

Finally, I have considered a model of monopolistic competition. In this model, even though the evolution of the industry reflects the choices made by all firms, each firm does not consider its decisions to have a material impact on the market. This is probably a reasonable assumption for the rather fragmented market for personal computers. However, in industries with fewer firms it may be desirable to account for strategic interaction in order to understand endogenous product differentiation.

2.8 Tables

Table 2.1: Descriptive statistics of lowest qualities sold by firms

Period	Mean	Std. Dev.	Min	Max
1	4.53	0.12	4.25	4.57
2	4.53	0.12	4.25	4.57
3	4.64	0.35	4.25	5.38
4	4.64	0.35	4.25	5.38
5	4.99	0.50	4.25	5.38
6	5.15	0.39	4.57	5.38
7	5.30	0.33	4.57	5.61
8	5.39	0.11	5.22	5.61
9	5.56	0.07	5.49	5.65
10	5.56	0.07	5.49	5.65
11	5.61	0.18	5.49	5.99
12	5.89	0.13	5.75	5.99
13	5.96	0.21	5.75	6.38
14	6.03	0.23	5.75	6.38
15	6.07	0.21	5.87	6.38
16	6.37	0.35	5.87	6.85
17	6.64	0.24	6.26	6.91
18	6.68	0.22	6.35	6.85
19	6.74	0.18	6.35	6.85
20	6.89	0.11	6.75	7.03
21	7.03	0.17	6.85	7.31
22	7.10	0.18	6.85	7.31
23	7.17	0.23	6.85	7.41
24	7.18	0.24	6.85	7.48
25	7.39	0.27	6.92	7.73
26	7.47	0.22	7.04	7.73

Table 2.2: Descriptive statistics of highest qualities sold by firms

Period	Mean	Std. Dev.	Min	Max
1	5.41	0.19	5.22	5.61
2	5.53	0.09	5.38	5.61
3	5.56	0.06	5.49	5.61
4	5.84	0.29	5.61	6.21
5	5.93	0.29	5.61	6.32
6	5.98	0.33	5.63	6.32
7	6.08	0.31	5.75	6.32
8	6.20	0.17	5.87	6.32
9	6.27	0.10	6.13	6.32
10	6.66	0.00	6.66	6.66
11	6.66	0.00	6.66	6.66
12	6.81	0.06	6.75	6.91
13	7.10	0.19	6.75	7.31
14	7.15	0.12	7.03	7.31
15	7.15	0.12	7.03	7.30
16	7.15	0.12	7.03	7.31
17	7.40	0.13	7.12	7.48
18	7.44	0.04	7.41	7.48
19	7.44	0.04	7.41	7.48
20	7.58	0.12	7.41	7.73
21	7.97	0.11	7.90	8.20
22	8.00	0.12	7.83	8.20
23	8.00	0.12	7.83	8.20
24	8.01	0.12	7.83	8.20
25	8.01	0.12	7.83	8.20
26	8.16	0.19	7.83	8.42

Table 2.3: Estimation of demand

Coefficient	OLS estimates (standard error)	OLS estimates (standard error)	IV estimates (standard error)	IV estimates (standard error)
Quality (α)	1.59* (0.18)	0.93* (0.19)	2.00* (0.22)	1.46* (0.46)
Price (γ)	−1.86* (0.13)	−1.49* (0.13)	−2.27* (0.18)	−2.08* (0.18)
RAM		−0.001 (0.002)		−0.0003 (0.002)
Hard disk		0.03* (0.01)		−0.03* (0.01)
CD − ROM		0.69* (0.22)		0.59* (0.22)
Modemspeed		0.02* (0.002)		0.02* (0.002)
Ethernet		−0.65* (0.15)		−0.59* (0.15)
DVD		0.45* (0.15)		0.46* (0.15)
Soundcard		0.07 (0.14)		0.08 (0.14)
Monitor		2.60* (0.36)		2.19* (0.37)
Monitorsize		−0.15* (0.02)		−0.14* (0.02)
R^2	0.46	0.53	0.45	0.53
N	1,703	1,703	1,703	1,703

* : statistically significant at the 1% level

Table 2.4: Regression of time-vendor dummies

Period	Coefficient	Standard deviation
t1	-13.57	0.48
t2	-14.31	0.48
t3	-14.39	0.48
t4	-14.62	0.48
t5	-14.75	0.48
t6	-15.10	0.48
t7	-14.80	0.48
t8	-15.58	0.48
t9	-15.90	0.48
t10	-16.75	0.48
t11	-17.20	0.48
t12	-17.33	0.48
t13	-17.95	0.48
t14	-18.52	0.48
t15	-18.29	0.48
t16	-18.51	0.48
t17	-19.07	0.48
t18	-19.48	0.48
t19	-19.49	0.48
t20	-19.95	0.48
t21	-21.36	0.48
t22	-21.17	0.48
t23	-21.67	0.48
t24	-21.92	0.48
t25	-22.34	0.48
t26	-23.10	0.48
$R^2 = 0.99$		
$N = 182$		

Table 2.5: Estimation of the marginal cost of quality

Time	Coefficient	Standard error
1	379.41	13.55
2	349.20	12.32
3	336.78	12.07
4	335.97	11.51
5	332.71	12.54
6	339.51	12.22
7	330.83	11.45
8	289.35	10.37
9	279.89	10.12
10	271.12	8.53
11	242.13	8.32
12	244.54	7.86
13	242.80	7.66
14	213.94	7.94
15	188.85	8.15
16	161.06	8.49
17	156.78	8.69
18	139.37	9.10
19	116.92	9.24
20	117.87	8.75
21	131.30	6.38
22	123.58	6.59
23	115.21	6.61
24	105.53	6.20
25	96.12	6.78
26	93.09	6.11
$R^2 = 0.91$		
$N = 1,703$		

Table 2.6: Evolution of the frontier of technology

Parameter	Estimate (standard error)
$\lambda_{\bar{Q}}^1$	2.569 (1.01)
$\lambda_{\bar{Q}}^2$	0.719 (0.11)
$\lambda_{\bar{Q}}^3$	−2.383 (1.0)
R^2	0.98
N	25

Table 2.7: Evolution of the marginal cost of quality

Parameter	Estimate (standard error)
λ_{ϕ}^1	0.243 (0.11)
λ_{ϕ}^2	0.718 (0.11)
λ_{ϕ}^3	−0.028 (0.12)
R^2	0.98
N	25

Table 2.8: Evolution of the aggregate index of market activity

Parameter	Estimate (standard error)
λ_{Ω}^1	4.174 (2.09)
λ_{Ω}^2	0.789 (0.12)
λ_{Ω}^3	−6.91 (3.81)
R^2	0.99
N	25

Table 2.9: Estimation of dynamic parameters

Parameter	Estimate ($m)
Fixed cost of production (F)	72.3
Sunk cost of withdrawing products ($\underline{\Gamma}$)	2.2
Sunk cost of introducing new products ($\underline{\Gamma}$)	4.1

Table 2.10: Parameters used for the model simulation

Parameter	Value
α	2.00
γ	−2.27
β	0.97
$\lambda^1_{\overline{Q}}$	2.569
$\lambda^2_{\overline{Q}}$	0.719
$\lambda^3_{\overline{Q}}$	−2.383
λ^1_{ϕ}	0.243
λ^2_{ϕ}	0.718
λ^3_{ϕ}	−0.028
F	$72.3m$
$\underline{\Gamma}$	$2.2m$
M	$1.8m$

Table 2.11: Initial conditions used for the model simulation

Variable	Initial value
$\overline{Q}$	5.408
$\underline{q}$	4.487
ϕ	0.3794

2.9 Figures

Figure 2-1: Evolution of firms' product lines

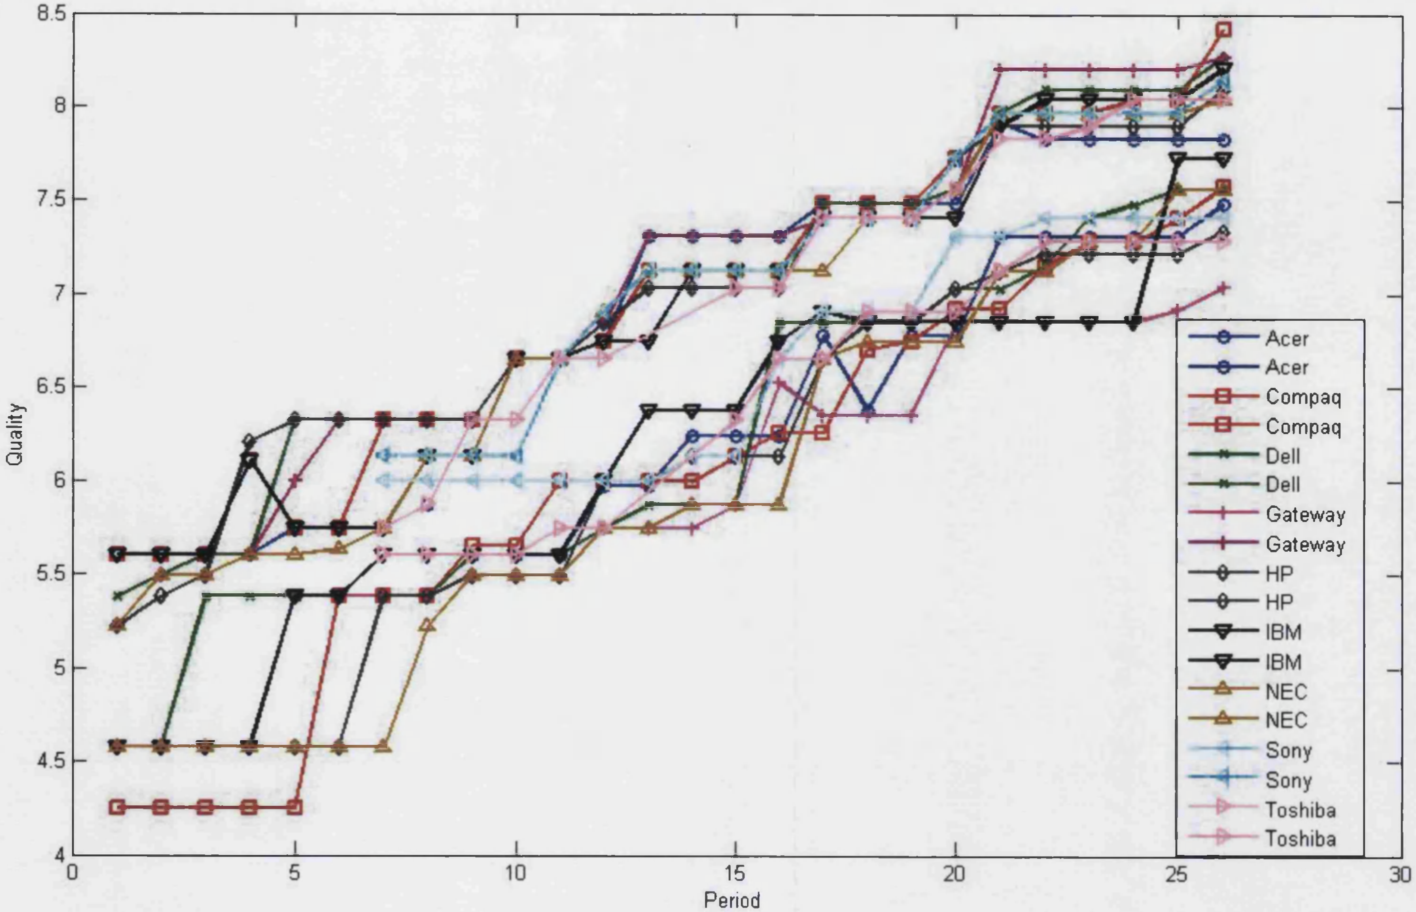

Figure 2-2: Evolution of firms' product lines

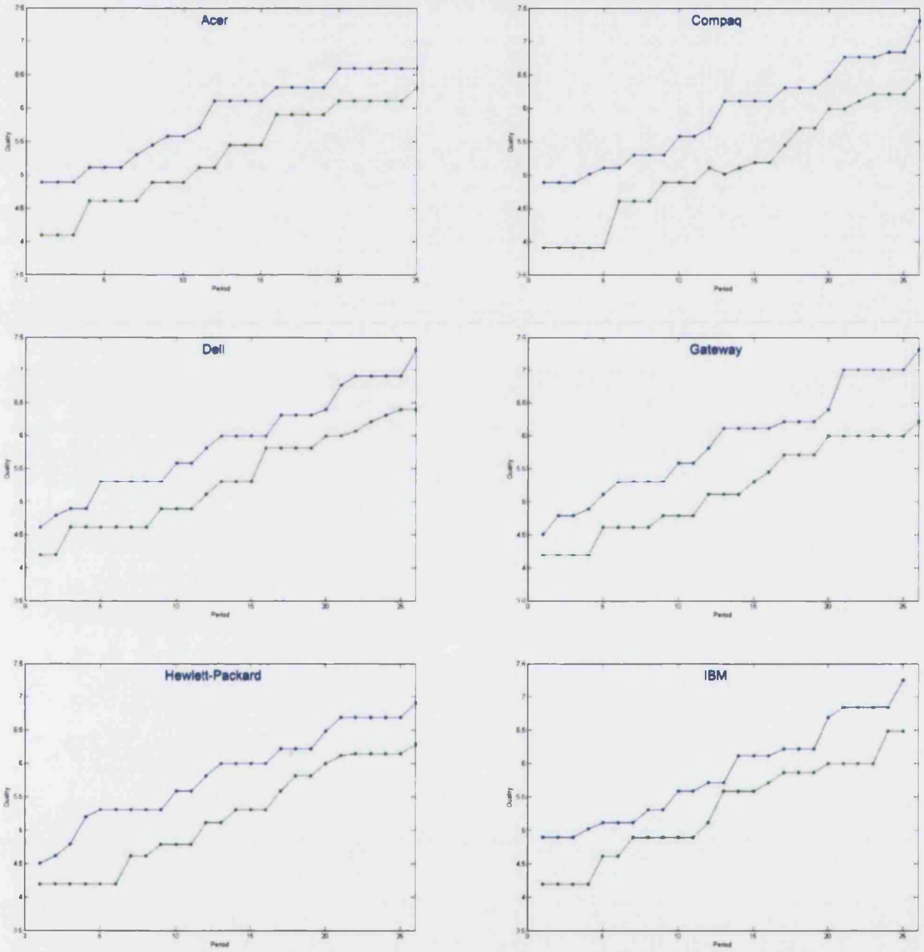

Figure 2-3: Simulated quality window

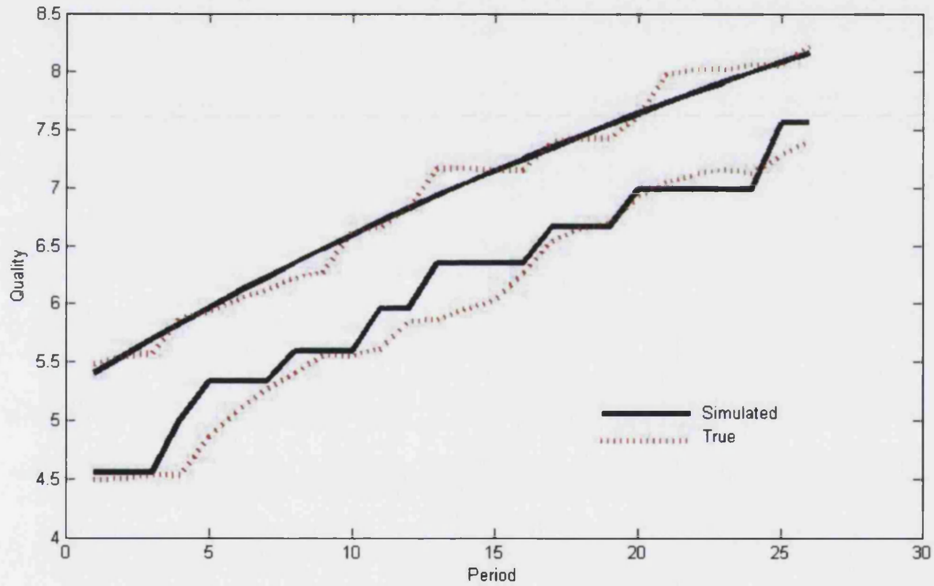

Figure 2-4: Simulation of evolution of marginal quality with no adjustment costs

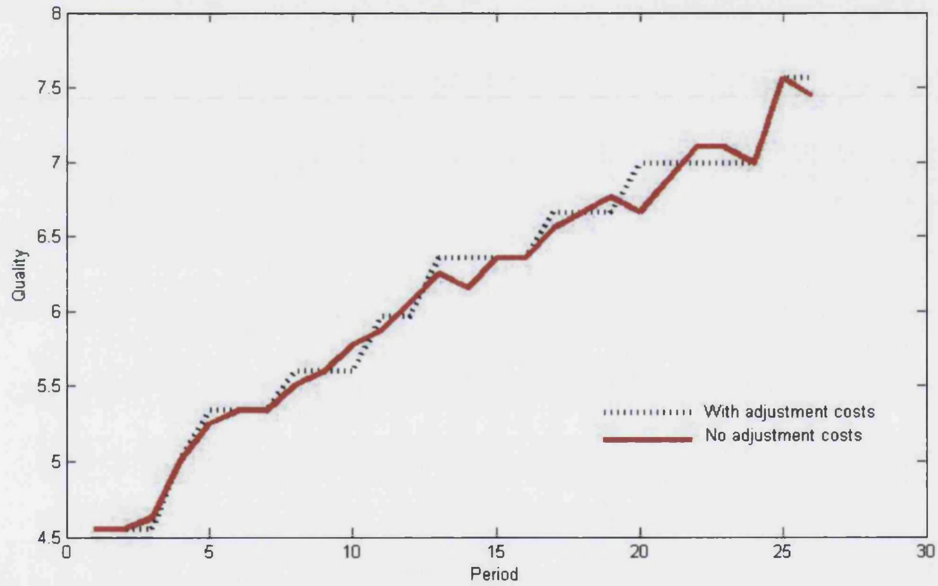

Chapter 3

Optimal product-line design, cannibalisation and product replacement

3.1 Introduction

In vertically-differentiated industries, technological progress often allows firms to sell new products of better quality. In some cases, innovation results in "product proliferation": the firm adds the new product to its product line alongside existing products thereby expanding the range of product-qualities supplied to consumers. For instance, in the market for laser printers, de Figueiredo and Kyle (2006) show that the introduction of new faster printers has not been associated with the exit of slower printers from the market.

Often, however, innovation has a destructive impact and the introduction of new better products is accompanied with the withdrawal of existing products from the market. In some cases, the destructive impact of innovation has a localised nature and it takes the form of product upgrading whereby an existing product is replaced by a new better version. For instance, it is common for car manufacturers to introduce new "generations" of their models with improved design and engineering, which often receive also a mid-life "facelift". In other markets, however, product innovation has a non-localised destructive impact that affects products that are distant in the characteristics space. For instance, as discussed in Chapters 1 and 2, the introduction of new better personal computers is associated with the exit of relatively obsolete personal computers at the bottom of the quality ladder.

In the previous Chapters, I developed two empirical structural models to study equilibrium product lines the market for personal computers and their evolution in the presence of ongoing technological progress. In these models, I explained the equilibrium range of qualities sold in the market as the result of competition among a set of monopolistically competitive firms. The analysis abstracted from strategic competition and from cannibalisation considerations but it nevertheless captured well the dynamic process of economic obsolescence in the market for personal computers in the form of a moving quality window driven by competitive forces.

In this Chapter I complement the analysis of the previous two Chapters by studying the cannibalisation considerations that shape a firm's vertically-differentiated optimal product line in the presence of product innovation. To be clear, the purpose of the analysis in this Chapter is not to investigate the market for personal computers further, but to develop a more general theoretical analysis of product exit decisions driven by cannibalisation considerations in vertically differentiated markets. By its nature,.the model that I discuss is extremely general and can be applied to understand product-line decisions observed in a variety of industries. In particular it is best suited to explain the localised product replacement decisions that are common in many dynamic markets such as the one for cars and the one for digital cameras and, by contrast, the proliferation of products in the quality space that is common in other industries such as the one for laser printers.

In order to identify under what circumstances innovation is accompanied with creative desctruction rather than with product proliferation, I consider a vertical-differentiation model under very general specifications of utility and technology in which a monopolist optimally chooses what products to sell among an exogenous set of feasible product-qualities. I focus my analysis on the behaviour of a monopolist in order to be able to study intra-firm cannibalisation considerations without having to worry about the effects of inter-firm competition on product-line decisions (as those that I investigated in the first two chapters of this thesis).

I study the impact that supplying a (new) product has on the incentives of the monopolist to include other products in the optimal product line and I present a necessary and sufficient condition for product replacement to be optimal. My model generalises the analysis of Siebert (2003) who concluded that an innovator is always better off by withdrawing the lower quality product in order to avoid cannibalisation. In particular, I show that product exit is optimal only in some cases (which I state as restrictions to the surplus function) and that Siebert's result stems from his specific functional specification of utility and technology.

The results that I obtain suggest also that cannibalisation considerations can well explain localised product replacement in the form of product upgrading that is witnessed in many markets (e.g. the market for cars and the market for digital cameras) where the introduction of a new product is associated with the exit of a neighbouring product in the quality space. However, cannibalisation considerations appear to be less relevant to explain the relationship between the introduction of new products and the exit of products that are distant in the quality space such as the one underlying the "moving quality window" observed in the market for personal computers that I studied in the previous Chapters, where shifts of the frontier of the technology are associated with the ordered exit of products at the bottom end of the quality ladder.

The next section briefly surveys the literature related to this study. Section 3.3 lays out the model. Section 3.4 presents the analysis of cannibalisation and optimal product line design. Final conclusions follow.

3.2 Related literature

From a modelling perspective, the vertical differentiation models that are most closely related to this study are those anlaysed by de Fraja (1996) and by Johnson and Myatt (2003, 2006), who also consider a specification of the vertical differentiation model in which there is an exogenous discrete set of feasible quality locations. de Fraja (1996) considers an oligopoly model of vertical differentiation with multiproduct firms competing à la Cournot and derives a number of results proving existence and uniqueness of equilibrium. Johnson and Myatt (2004,2006) consider also a model in which multi-product firms choose the optimal set of qualities to supply. They develop an original approach to analyse the model, based on the notion of "upgrades", and present a number of interesting results concerning equilibrium product lines and the impact of competition of the products sold in the market. Even though I do not consider an oligopolistic model, I study optimal product lines under a more general specification of utility and technology than de Fraja's and my analysis does not follow Johnson and Myatt's upgrade approach.

Given my focus on product replacement, this study also contributes to the literature on product exit, which spans the boundary between economics and marketing.

Stavins (1995) considers entry and exit at the model level in the market for personal computers. She estimates a logit model of the probability of exit in a given year, in which explanatory variables include measures of overpricing of models as well as firm and model characteristics. She concludes that overpriced models are more likely to exit the market. The probability of exit is also higher for products that have been on the market for longer and for products sold by new entrants.

Greenstein and Wade (1998) investigate product life cycles in the commercial mainframe computer market. They also find that as products age their probability of exiting the market increases. They also show that exit rates are influenced by competition from within a product's size class as well as from competition in surrounding niches. In addition, they also find that once a firm introduces a new product in a size class, other products from that firm in that niche are more likely to exit.

De Figueiredo and Kyle (2006) study product turnover in the laser print industry. They show that a product is more prone to exit the older it is and the more products are in the market and in the same niche. Printers on the frontier of resolution survive longer on the market than less advanced products. Firms with strong brands are less likely to withdraw their products from the market. Contrary to the findings of Greenstein and Wade they do not find that exit is more likely when a firm introduces an additional product in the same niche.

Requena and Walker (2003) study product entry and exit in the car industry and show that exit is mainly driven by the introduction of new models in the same market segment. Genakos (2005) documents instead the shortening of product life cycles in the market for personal computers. Chisholm and Norman (2005) study optimal exit decisions in the U.S. motion-pictures exhibition market.

The empirical literature, however, often laments the lack of theoretical foundations to the study of product exit. Indeed, there are only few theoretical papers that consider the decision by a firm to withdraw existing products.

Judd (1985) considers the decision by an incumbent to keep or withdraw a product in the presence of entry in a horizontally differentiated market. One the one hand, keeping the product in the market increases sales but reduces the price of its other products. On the other hand, withdrawing the product reduces sales but increase the prices of its other products. In Judd's horizontal differentiation model, the firm is better off by withdrawing the products that are close to those sold by the entrant to soften price competition.

Siebert (2003) considers product withdrawal in a vertical-differentiation model and concludes that the innovator is always better off by withdrawing the lower quality product in order to avoid cannibalisation. By considering a more general specification of the model, I am able to show that it is not always optimal for the monopolist to withdraw the low-quality product from the market and, in fact, I derive a necessary and sufficient condition for product replacement to be optimal.

3.3 The model

There is a population of heterogeneous consumers, whose willingness to pay for quality θ is distributed according to the cumulative distribution function $F(\theta)$ with support on $\left[\underline{\theta}, \overline{\theta}\right] \in \Re^{+}$. I assume that $F(\theta)$ is continuous and twice differentiable.

A consumer of type θ derives net utility $u(\theta, q_i) - p_i$ from the consumption of good of quality q_i purchased at price p_i. I assume that $u(\theta, q') - u(\theta, q)$ is strictly increasing in θ whenever $q' > q$, i.e. that $u_\theta(\theta, q') > u_\theta(\theta, q)$. This condition implies that consumers of higher type value an increase in the quality of the good more than consumers of lower type and ensures that the indifference curves of any two types of consumers cross at most once in price-quality space and that the associated demand curves do not intersect.

The monopolist can supply any of N feasible qualities in the exogenous set $\Omega = \{q_1, q_2, ..., q_N\}$. Quality q_i can be produced at constant marginal cost $c(q_i) = c_i$. Denote the quality of the outside option by q_0 and normalise its price p_0 to zero.

The monopolist offers a menu of feasible qualities $(q \in \Omega, p(q))$ to maximise expected profits, given that consumers will select from the schedule to maximise individual utilities. Following the revelation approach, I consider the mechanism

$$M = \{p(\theta), q(\theta) \in \Omega, \theta \in [\underline{\theta}, \overline{\theta}]\}$$

whereby a consumer who reports type θ is assigned quality $q(\theta)$, belonging to the set of feasible qualities Ω, and pays price $p(\theta)$.

According to the revelation principle, in order to solve for the optimal mechanism, it is possible to focus on those allocations that can be truthfully implemented.

An allocation can be truthfully implemented if it is individually rational and incentive compatible. Individual rationality requires that any consumer who purchases a good $q(\theta)$ at price $p(\theta)$ enjoys a utility that is no lower than that of the outside option:

$$u(\theta, q(\theta)) - p(\theta) \geq u(\theta, q_0)$$

Incentive compatibility requires that, conditional on buying some good, a consumer of type θ_i must (weakly) prefer allocation $(q(\theta_i), p(\theta_i))$ rather than any other allocation offered by the monopolist:

$$u(\theta_i, q(\theta_i)) - p(\theta_i) \geq u(\theta_i, q(\theta_j)) - p(\theta_j)$$

for any $\theta_j \neq \theta_i$.

An equilibrium allocation is an implementable allocation that maximises the monopolist's profits. The optimal allocation therefore solves:

$$\max_{q(\theta)\in\Omega, p(\theta)} \int_{\underline{\theta}}^{\overline{\theta}} [p(\theta) - c(q(\theta))]\, dF(\theta)$$

subject to

$$u(\theta, q(\theta)) - p(\theta) \geq u(\theta, q(\theta')) - p(\theta') \text{ for any } \theta,\ \theta' \neq \theta$$

and

$$u(\theta, q(\theta)) - p(\theta) \geq u(\theta, q_0)$$

I now present a number of results that allow me to simplify the maximisation problem that determines the equilibrium allocation. First of all, I present two standard results on implementable allocations. I then suggest that, in any implementable allocation, there is a one-to-one relationship between the quality allocation $q(\theta)$ and a vector of marginal consumers $\boldsymbol{\theta} = [\theta_1, \theta_2, .., \theta_N]$. This implies that choosing $q(\theta)$ is equivalent to choosing N marginal consumers θ_i such that all consumer types $\theta : \theta_{i+1} > \theta \geq \theta_i$ are allocated, i.e. consume, quality q_i. I then identify the unique price schedule $\mathbf{p}$ that optimally implements a given quality allocation $\boldsymbol{\theta}$. These results will allow me to rewrite the maximisation problem of the monopolist as the choice of N marginal consumers, i.e. the vector $\boldsymbol{\theta}$, given the optimal implementing price schedule $\mathbf{p}$.

Claim 3.1 *(i) In any implementable allocation, all consumers who are allocated the same quality must pay the same price, i.e.* $p(\theta) = p(q(\theta))$*.(ii) In any implementable allocation, higher-type consumers must purchase products of weakly higher quality, i.e. if* $\theta_j > \theta_i$ *then* $q(\theta_j) \geq q(\theta_i)$*.*

Proof. See Appendix 3.6.1. ∎

The above (standard) result states that any implementable allocation, and therefore also the equilibrium allocation, must be weakly monotonic: consumers of higher type purchase goods of weakly higher quality.

Given this result, in any implementable allocation, let $\theta_i \in [\underline{\theta}, \overline{\theta}]$, $i = 1, 2, ..., N$, be the "marginal" consumer type such that all types $\theta : \theta_{i+1} \geq \theta \geq \theta_i$ are allocated quality q_i. Clearly, weak monotonicity implies that $\theta_j \geq \theta_i$ for any $j > i$. Notice that the vector $\boldsymbol{\theta}$ also pins down the supply vector $\mathbf{x} = [x_1, x_2, ..., x_N]$ where $x_i = F(\theta_{i+1}) - F(\theta_i) \geq 0$ is the supply of good of quality q_i. Therefore, quality q_i is sold if and only if $\theta_{i+1} > \theta_i$.

Lemma 3.2 states that there is a one-to-one relationship between any implementable quality allocation $q(\theta)$ and the vector $\boldsymbol{\theta} \equiv [\theta_1, \theta_2, ..., \theta_N]$ of N marginal consumers θ_i.

Lemma 3.2 *Any implementable quality allocation* $q(\theta)$ *can be described by a vector of marginal consumers* $\boldsymbol{\theta} = [\theta_1, \theta_2, .., \theta_N]$*. Any vector* $\boldsymbol{\theta}$ *describes a unique weakly monotonic schedule* $q(\theta)$*.*

Proof. Straightforward given the weak monotonicity of implementable allocations and the definition of $\boldsymbol{\theta}$. ∎

In order to get rid of incentive compatibility and rationality constraints in the monopolist's maximisation problem a further step is needed. For any quality allocation $\boldsymbol{\theta}$, I identify the optimal implementing price vector $\mathbf{p}(\boldsymbol{\theta})$, i.e. the price vector that maximises the monopolist's profits conditional on the resulting allocation being implementable.

Lemma 3.3 *The price vector* $\mathbf{p}(\boldsymbol{\theta}) = [p_1, p_2, ..., p_N]$ *that optimally implements quality allocation* $\boldsymbol{\theta}$ *has element*

$$p_i = \sum_{l=1}^{i} u(\theta_l, q_l) - u(\theta_l, q_{l-1}) \tag{3.1}$$

Proof. See Appendix 3.6.2. ∎

Notice that expression 3.1 applies also if for some i $\theta_i = \theta_{i-1}$, i.e. if some feasible qualities are not sold by the monopolist.

With these results in hand, I can now rewrite and solve the simplified optimisation problem of the monopolist. Given the monotonicity property of any implementable quality allocation, focusing the search of the optimal allocation within the set of weakly monotonic allocations where $\theta_j \geq \theta_i$ for any $j > i$, I can write the profit of the monopolist as the sum of profits from each exogenous feasible quality $q_1, q_2, ..., q_N$ as follows:

$$\begin{aligned}
\Pi &= \int_{\theta_1}^{\theta_2} [p(q_1) - c(q_1)]\, dF(\theta) + \int_{\theta_2}^{\theta_3} [p(q_2) - c(q_2)]\, dF(\theta) + ... + \int_{\theta_N}^{\overline{\theta}} [p(q_N) - c(q_N)]\, dF(\theta) \\
&= [p(q_1) - c(q_1)]\,[F(\theta_2) - F(\theta_1)] + ... + [p(q_N) - c(q_N)]\,\left[F\left(\overline{\theta}\right) - F(\theta_N)\right] \\
&= \sum_{i=1,2,...,N} [p(q_i) - c(q_i)]\,[F(\theta_{i+1}) - F(\theta_i)]
\end{aligned}$$

where I define $\theta_{N+1} \equiv \overline{\theta}$.

I can now substitute the optimal implementing prices

$$p_i = \sum_{l=1}^{i} u(\theta_l, q_l) - u(\theta_l, q_{l-1})$$

in the expression for profits and consider the following maximisation problem in which the monopolist solves for the optimal vector $\boldsymbol{\theta}^* = [\theta_1^*, \theta_2^*, ..., \theta_N^*]$ under the monotonicity constraints $\theta_j \geq \theta_i$ for any $j > i$:

$$\max_{\boldsymbol{\theta}} \Pi(\boldsymbol{\theta}) = \sum_{l=1,2,...,N} \left[\sum_{m=1}^{l} \left(u(\theta_m, q_m) - u(\theta_m, q_{m-1})\right) - c_l \right] \left[F(\theta_{l+1}) - F(\theta_l)\right] \tag{3.2}$$

subject to

$$\theta_i - \theta_j \leq 0 \text{ for any } i < j$$

The Lagrangian of this maximisation problem is

$$L(\boldsymbol{\theta}) = \Pi(\boldsymbol{\theta}) + \sum_{l=1}^{N-1} \sum_{m=i+1}^{N} \left[\lambda_{lm}(\theta_l - \theta_m)\right] \tag{3.3}$$

The Kuhn and Tucker's necessary conditions for maximisation are:[1]

$$\frac{\partial L(\boldsymbol{\theta}^*)}{\partial \theta_i} = \frac{\partial \Pi(\boldsymbol{\theta}^*)}{\partial \theta_i} - \sum_{h=1}^{i-1} \lambda_{hi}^* + \sum_{h=i+1}^{N} \lambda_{ih}^* = 0 \tag{3.4}$$

$$\lambda_{ij}^* \left(\theta_i^* - \theta_j^*\right) = 0$$

$$\theta_i^* - \theta_j^* \leq 0$$

[1] If $\frac{\partial^2 \Pi}{\partial \theta_i^2} < 0$ for any θ_i, given that $\frac{\partial^2 \Pi}{\partial \theta_i \partial \theta_j} = 0$ for any $i \neq j$, then the Hessian matrix is negative semidefinite and the necessary Kuhn and Tucker's conditions for a maximum are also sufficient.

$$\lambda_{ij}^* \leq 0 \tag{3.5}$$

In Appendix 3.6.3 I show that

$$\begin{aligned}\frac{\partial \Pi(\boldsymbol{\theta}^*)}{\partial \theta_i} &= -\left[u\left(\theta_i^*, q_i\right) - u\left(\theta_i^*, q_{i-1}\right) - \left(c_i - c_{i-1}\right)\right] f\left(\theta_i^*\right) + \\ &\quad + \left[1 - F\left(\theta_i^*\right)\right]\left[u_\theta\left(\theta_i^*, q_i\right) - u_\theta\left(\theta_i^*, q_{i-1}\right)\right]\end{aligned} \tag{3.6}$$

where $c_0 = 0$.

Define:

- $\Delta S(\theta, q_j, q_i) \equiv [u(\theta, q_j) - c_j] - [u(\theta, q_i) - c_i]$, i.e. the change in social surplus associated to consumption of quality q_j rather than quality q_i by a consumer of type θ ("upgrade surplus function"); and
- $\Delta S_\theta(\theta, q_j, q_i) \equiv \frac{\partial}{\partial \theta} \Delta S(\theta, q_j, q_i) = [u_\theta(\theta, q_j) - u_\theta(\theta, q_i)]$, i.e. the derivative of the upgrade surplus function with respect to the consumer-type θ. Notice that $\Delta S_\theta(\theta, q_j, q_i) > 0$ for any $q_j > q_i$ because, by assumption, $[u(\theta, q_j) - u(\theta, q_i)]$ is increasing in θ for any $q_j > q_i$.

Expression 3.6 can thus be rewritten as

$$\frac{\partial \Pi(\boldsymbol{\theta}^*)}{\partial \theta_i} = -\Delta S\left(\theta_i^*, q_i, q_{i-1}\right) f\left(\theta_i^*\right) + \left[1 - F\left(\theta_i^*\right)\right] \Delta S_\theta\left(\theta_i^*, q_i, q_{i-1}\right) \tag{3.7}$$

3.4 Analysis: cannibalisation and product replacement

In this section I use the theoretical model to study the conditions under which different feasible product-qualities can coexist in the optimal product line. This is useful to understand when the introduction of a new product in the product line is associated with product proliferation and when it results instead in product replacement.

First of all, I discuss if a product can ever have a destructive impact on products of higher quality because of cannibalisation considerations. I then consider the destructive impact that introducing a product in the product line may have on lower-quality products: I show that product replacement due to cannibalisation considerations is a localised phenomenon and I present a necessary and sufficient condition for product replacement (rather than product proliferation) to be optimal.

The first question that I address is if the introduction of a product can ever have a destructive impact on products of higher quality, i.e. if selling the high-quality product cannibalises sales of the low-quality product that ,may be optimal to avoid by not selling the high-quality product. In order to more clearly discuss the relevant economic mechanism, it is useful to focus on the simplified case in which the set of feasible qualities includes only two qualities, q_j and q_i with $q_j > q_i$.

Proposition 3.4 states that, as long as there are some consumers who prefer to purchase the higher-quality good rather than the lower-quality good if both were to be sold at prices equal to their respective marginal costs, then it is always optimal to sell the good of higher quality alongside the good of lower quality (or possibly alone).

Proposition 3.4 *If the monopolist sells quality q_i, then selling quality q_j, $q_j > q_i$, is not optimal if and only if $\Delta S(\theta, q_j, q_i) < 0$ for every θ.*

Proof. If the firm sells only quality q_i, to consumers $\theta \geq \theta_i$, it obtains profits equal to

$$\pi(\theta_i) = [1 - F(\theta_i)][u(\theta_i, q_i) - u(\theta_i, q_0) - c_i]$$

If the firm sells quality q_j alongside quality q_i to all consumer types $\theta \geq \theta_j > \theta_i$ it obtains profits

$$\begin{aligned}\widetilde{\pi}(\theta_i, \theta_j) &= [1 - F(\theta_i)][u(\theta_i, q_i) - u(\theta_i, q_0) - c_i] + \\ &+ [F(\theta_j) - F(\theta_i)][u(\theta_j, q_j) - u(\theta_j, q_i) - (c_j - c_i)]\end{aligned}$$

which can be written as

$$\widetilde{\pi}(\theta_i, \theta_j) = \pi(\theta_i) + [F(\theta_j) - F(\theta_i)] \Delta S(\theta_j, q_j, q_i)$$

To prove sufficiency, notice that, if $\Delta S(\theta, q_j, q_i) < 0$ for every θ, then there is no $\theta_j > \theta_i$ such that $\widetilde{\pi}(\theta_i, \theta_j) > \pi(\theta_i)$.

To prove necessity, suppose that there is a consumer-type θ^* such that $\Delta S(\theta^*, q_j, q_i) > 0$ and that $\Delta S(\theta, q_j, q_i) < 0$ for every $\theta < \theta^*$. Since $\Delta S_\theta(\theta, q_j, q_i) > 0$, $\Delta S(\theta, q_j, q_i) > 0$ for every $\theta > \theta^*$. In this case, the firm does better than selling only quality q_i by selling quality q_j to all $\theta \geq \theta_j^*$ where $\theta_j^* \equiv \max(\theta_i, \theta^*)$.

■

A product of higher quality cannot be optimally sold alongside a product of lower quality only if all consumers prefer to purchase the latter rather then the former if both were sold at a prices equal to their respective marginal costs of production, i.e. there is no consumer who is willing to pay a premium for the high quality good that is at least as large as the incremental cost of production.

Selling quality q_j alongside quality q_i does not increase overall sales but it reduces the sales of quality q_i. This cannibalisation, however, is only apparent: when a unit of the high quality good is sold, the consumer is implicitly purchasing a unit of the low quality good in addition to an "upgrade". This can be seen by considering the firm's profit function

$$\begin{aligned}\widetilde{\pi} &= \left[F\left(\theta_j\right) - F\left(\theta_i\right)\right]\left[u\left(\theta_i, q_i\right) - u\left(\theta_i, q_0\right) - c_i\right] + \\ &+ \left[1 - F\left(\theta_j\right)\right]\left[u\left(\theta_j, q_j\right) - u\left(\theta_j, q_i\right) + u\left(\theta_i, q_i\right) - u\left(\theta_i, q_0\right) - c_j\right]\end{aligned}$$

which shows that the firm sells quality q_j to $\left[1 - F\left(\theta_j\right)\right]$ consumers and quality q_i to $\left[F\left(\theta_j\right) - F\left(\theta_i\right)\right]$ consumers and by observing that it can be rewritten as

$$\begin{aligned}\widetilde{\pi} &= \left[1 - F\left(\theta_i\right)\right]\left[u\left(\theta_i, q_i\right) - u\left(\theta_i, q_0\right) - c_i\right] + \\ &+ \left[F\left(\theta_j\right) - F\left(\theta_i\right)\right]\left[u\left(\theta_j, q_j\right) - u\left(\theta_j, q_i\right) - \left(c_j - c_i\right)\right]\end{aligned}$$

The expression shows that the firm can be thought of as selling quality q_i to $\left[1 - F\left(\theta_i\right)\right]$ consumers and the upgrade $\left[q_j - q_i\right]$ (which costs $\left(c_j - c_i\right)$ to produce) to $\left[F\left(\theta_j\right) - F\left(\theta_i\right)\right]$ consumers.[2] If no consumer is willing to pay for the upgrade a price that is at least equal to the cost of the upgrade, then, as proposition 3.4 shows, it is not optimal to supply the upgrade, i.e. the high quality product. Therefore, the reason why it might not be profitable for the monopolist to supply a high-quality product alongside the low-quality product is that consumers do not value the "upgrade" enough to cover the higher cost of production of the high-quality good and not the fact that the high-quality good cannibalises the low-quality good.

In this (standard) vertical differentiation model, it is the low-quality good that "cannibalises" the high-quality good in the sense the former constrains the price at which the latter can be sold. Product replacement driven by cannibalisation considerations therefore has a downward nature: it is the in-

[2]Johnson and Myatt (2003, 2006) have used this notion of "upgrades" to develop the analysis of a Cournot vertical-differentiation oligopoly.

troduction of a product of higher quality in the product line that may cause the firm to optimally withdraw products of lower quality from the market in order to increase the price at which the (new) high quality good can be sold.

However, in the presence of multiple feasible qualities, it is not clear what low-quality products are affected by this potential destructive effect: does it affect products at the bottom of the quality ladder or is it relevant only "locally"?

Proposition 3.5 states that product exit stemming from cannibalisation considerations has always a localised nature in the sense that if q_i is not optimally sold alongside quality q_j, with $q_j > q_i$, then no other intermediate quality $q_m : q_i \leq q_m < q_j$ can optimally be sold alongside quality q_j.

Proposition 3.5 *Consider three qualities q_j, q_m and q_i, with $q_j > q_m > q_i$. If $\theta_j^* = \theta_i^*$, then $\theta_j^* = \theta_m^*$.*

Proof. The statement follows directly from the weak monotonicity constraint of implementable allocations. If $\theta_j^* = \theta_i^*$ but $\theta_j^* > \theta_m^*$ then it would follow that $\theta_i^* > \theta_m^*$ which contradicts the requirement that $\theta_j^* \geq \theta_i^*$ for any $q_j > q_i$. ■

This result does not depend on any restriction on the utility function and the cost function other that the standard assumption that consumers of higher type value an increase in the quality of the good more than consumers of lower type. In fact, it follows directly from the localised nature of the vertical differentiation model. However, it has an interesting implication for the study of product replacement: proposition 3.5 suggests that cannibalisation considerations can explain product replacement, but only in the form of localised product upgrading observed in markets such as that for cars in which firms "upgrade" their products. By contrast, cannibalisation considerations in this model do not appear relevant to explaining product replacement as observed in a market like that for personal computers, where movements of the frontier of technology are associated with the relatively ordered exit of products at the bottom of the quality ladder.

Having determined that product replacement driven by cannibalisation considerations is a localised phenomenon that affects neighbouring qualities, the following proposition provides a necessary and sufficient condition for product replacement, rather than product proliferation, to be optimal.

Proposition 3.6 *Consider two qualities q_j and q_i, with $q_j > q_i$. Quality q_i cannot be optimally sold alongside quality q_j, i.e. $\theta_j^* = \theta_i^*$, if and only if*

$$\frac{\partial}{\partial \theta}\left(\frac{\Delta S\left(\theta, q_j, q_i\right)}{\Delta S\left(\theta, q_i, q_{i-}\right)}\right) \leq 0$$

where q_{i-} is the highest quality below q_i optimally sold by the monopolist.

Proof. See Appendix 3.6.4. ∎

Proposition 3.6 states that selling a product of quality q_i alognside a product of higher quality q_j cannot be optimal if and only if for higher-type consumers the upgrade surplus from quality q_i to quality q_j is proportionally lower than the upgrade surplus from quality q_{i-} to quality q_i. In other words, product replacement occurs if higher consumer types value the upgrade from good q_i to good q_j *proportionally* less than the upgrade from q_{i-} to q_i when all products are sold at prices equal to their respective marginal costs.

This general result can be applied to specific utility and production functions. I have considered two specific cases that are common in the literature: the case in which consumers have utility function $u(\theta, q) = \theta q$, and the case in which all qualities can be produced at the same marginal cost.

Define $\Delta q_j \equiv q_j - q_i$, $\Delta q_i \equiv q_i - q_{i-}$, $\Delta c_j = c_j - c_i$ and $\Delta c_i = c_i - c_{i-}$.

Corollary 3.7 *Suppose consumers' utility function is $u(\theta, q) = \theta q$. Quality q_i cannot be optimally sold alongside quality q_j, i.e. $\theta_j^* = \theta_i^*$, if and only if $\frac{\Delta c_j}{\Delta q_j} \leq \frac{\Delta c_i}{\Delta q_i}$.*

Proof. According to proposition 3.6, $\theta_j^* = \theta_i^*$ if and only if $\frac{\Delta S(\theta, q_j, q_i)}{\Delta S(\theta, q_i, q_{i-})}$ is decreasing in θ. If $u(\theta, q) = \theta q$, $S(\theta, q_j, q_i) = \theta \Delta q_j - \Delta c_j$ and therefore, a necessary and sufficient condition for $\frac{\partial}{\partial \theta}\left(\frac{\Delta S(\theta, q_j, q_i)}{\Delta S(\theta, q_i, q_{i-})}\right) \leq 0$ is that

$$\begin{aligned}
\Delta S_\theta(\theta, q_j, q_i) \Delta S(\theta, q_i, q_{i-}) - \Delta S_\theta(\theta, q_i, q_{i-}) \Delta S(\theta, q_j, q_i) &\leq 0 \rightarrow \\
\Delta q_j [\theta \Delta q_i - \Delta c_i] - \Delta q_i [\theta \Delta q_j - \Delta c_j] &\leq 0 \rightarrow \\
\theta \Delta q_i \Delta q_j - \Delta q_j \Delta c_i - \theta \Delta q_i \Delta q_j + \Delta q_i \Delta c_j &\leq 0 \rightarrow \\
\Delta q_i \Delta c_j - \Delta q_j \Delta c_i &\leq 0 \rightarrow \\
\frac{\Delta c_j}{\Delta q_j} &\leq \frac{\Delta c_i}{\Delta q_i}
\end{aligned}$$

■

Corollary 3.7 implies that, when the utility function is linear, product replacement occurs when the cost per unit of quality of the upgrade from q_i to q_j is lower than the cost per unit of quality of the upgrade from q_{i-} to q_i. This result is consistent with the one obtained by Johnson and Myatt (2003) who suggest that firms sell all feasible qualities if the utility function is $u(\theta, q) = \theta q$ and there are decreasing returns to quality, i.e. $\left(\frac{c_i - c_{i-1}}{q_i - q_{i-1}}\right)$ is increasing in i.

Corollary 3.8 *Suppose qualities q_j, q_i and q_{i-} can all be produced at the same marginal cost c. Quality q_i cannot be optimally sold alongside quality q_j, i.e. $\theta_j^* = \theta_i^*$, if and only if $\frac{\partial}{\partial \theta}\left(\frac{\Delta u(\theta, q_j, q_i)}{\Delta u(\theta, q_i, q_{i-})}\right) \leq 0$.*

Proof. If q_j, q_i and q_{i-} can all be produced at the same marginal cost c and $\frac{\partial}{\partial \theta}\left(\frac{\Delta u(\theta, q_j, q_i)}{\Delta u(\theta, q_i, q_{i-})}\right) \leq 0$, then $\frac{\Delta S(\theta, q_j, q_i)}{\Delta S(\theta, q_i, q_{i-})} = \frac{\Delta u(\theta, q_j, q_i)}{\Delta u(\theta, q_i, q_{i-})} \leq 0$. Proposition 3.6

states that this is a necessary and sufficient condition for product replacement to occur, i.e. $\theta_j^* = \theta_i^*$. ■

When an improvement in quality can be obtained as the result of R&D investments without a significant impact on the cost of production (e.g. software), the results obtained suggest that product replacement takes place when consumers of higher type value the quality upgrade proportionally less than low type consumers.

Proposition 3.6 is also useful to explain the extreme result obtained by Siebert (2003), according to which product replacement is always optimal. In fact, Siebert considers a model in which consumers have a linear utility function and all qualities can be produced at the same marginal cost. The combination of these assumptions implies that $\frac{\partial}{\partial\theta}\left(\frac{\Delta S(\theta,q_j,q_i)}{\Delta S(\theta,q_i,q_{i-})}\right) = 0$ and proposition 3.6 suggests that in this case product proliferation cannot be optimal.

3.5 Conclusions

In this Chapter I have presented a model of vertical differentiation to study the impact that product innovation has on the survival of other products in a firm's product line because of cannibalisation considerations. In general, a (new) better product may coexist with other products sold by the firm (product proliferation) or may not (product replacement). The analysis presented in this Chapter suggests that, in a vertically-differentiated environment, product replacement stemming from cannibalisation considerations affects the survival of products of lower quality and that it has a localised nature, i.e. it takes the form of "product upgrading" whereby the new product replaces neighbouring products in the quality ladder. I have derived a necessary and sufficient condition for product innovation to be associated with product proliferation rather than with product replacement and I have considered how this condition applies to some common parametric specifications of the vertical differentiation model.

The analysis can be extended in various directions. In particular, it might be interesting to study how endogenous product replacement or product proliferation affects the incentives of a multiproduct monopolist to invest in innovation, especially in a dynamic perspective in which the firm anticipates the impact of its future innovative activity on the profitability of the products that it develops.

3.6 Appendix to Chapter 3

3.6.1 Proof of claim 3.1

Proof. (i) Suppose that for two different consumer types $\theta_i \neq \theta_j$ the allocation is such that $q(\theta_i) = q(\theta_j)$ but $p(\theta_i) > p(\theta_j)$. In this case, a type θ_i consumer would enjoy a higher utility by reporting θ_j and consuming $(q(\theta_j) = q(\theta_i), p(\theta_j))$, i.e. it would not have the incentive to truthfully report its true type. The allocation considered cannot therefore be truthfully implemented since the incentive compatibility constraint would be violated.

(ii) I will prove this claim by contradiction. Consider an allocation such that a consumer of type θ_i purchases product q_i at price p_i and a consumer of type $\theta_j > \theta_i$ purchases product q_j at price p_j with $q_j < q_i$, i.e. the higher type consumes a product of lower quality.

If the proposed allocation were implementable it would have to be incentive compatible, i.e. a consumer of type θ_i would have to derive more utility from purchasing product (q_i, p_i) rather than product (q_j, p_j):

$$u(\theta_i, q_i) - p_i \geq u(\theta_i, q_j) - p_j$$

which can be written as

$$u(\theta_i, q_i) - u(\theta_i, q_j) \geq p_i - p_j \tag{3.8}$$

Similarly, implementability requires that a consumer of type θ_j must derive higher utility from purchasing product (q_j, p_j) rather than product (q_i, p_i):

$$u(\theta_j, q_j) - p_j \geq u(\theta_j, q_i) - p_i$$

which can be written as

$$p_i - p_j \geq u(\theta_j, q_i) - u(\theta_j, q_j) \tag{3.9}$$

Expressions (1) and (2) imply that the given allocation is implementable only if

$$u(\theta_i, q_i) - u(\theta_i, q_j) \geq u(\theta_j, q_i) - u(\theta_j, q_j)$$

However, this contradicts the assumption whereby $u(\theta, q') - u(\theta, q)$ is increasing in θ when $q' > q$. The proposed allocation cannot therefore be truthfully implemented. ■

3.6.2 Proof of claim 3.3

Proof. A price vector $\mathbf{p}$ optimally implements the quality allocation described by the vector $\boldsymbol{\theta}$ if the resulting allocation satisfies all incentive compatibility and rationality constraints and if the monopolist achieves the highest profits from quality allocation $\boldsymbol{\theta}$. The proof shows that only prices $p_i = \sum_{l=1}^{i} u(\theta_l, q_l) - u(\theta_l, q_{l-1})$ satisfy these requirements.

I prove the statement by induction. Consider the lowest quality offered by the monopolist q_1. By definition, $\theta_1 \geq \underline{\theta}$ is the consumer type such that all consumers $\theta \geq \theta_1$ purchase a good of quality q_1 or higher.

If the monopolist charges a price p_1^+ higher than $p_1 = u(\theta_1, q_1) - u(\theta_1, q_0)$ the quality allocation described by the vector $\boldsymbol{\theta}$ cannot be implemented since all consumers $\theta_1^+ < \theta \leq \theta_1$ where θ_1^+ solves $p_1^+ = u\left(\theta_1^+, q_1\right) - u\left(\theta_1^+, q_0\right)$ would be better off by purchasing the outside option q_0 rather than good of quality q_1.

On the other hand, it would not be optimal for the monopolist to charge a lower price since by increasing the price up to $p_1 = u(\theta_1, q_1) - u(\theta_1, q_0)$, he could still implement the quality allocation $\boldsymbol{\theta}$ and obtain higher profits

Similarly, consider now the second lowest quality q_2. In an implementable allocation in which consumers of type θ, $\theta \geq \theta_2$, purchase good of quality q_2 or higher, the monopolist cannot charge a price higher than $p_2 = \sum_{l=1}^{2} u(\theta_l, q_l) - u(\theta_l, q_{l-1})$. If it did so, some consumers $\theta \geq \theta_2$ would be better off purchasing good of quality q_1 at price p_1. On the other hand, the monopolist has no incentive to charge a price lower than $p_2 = \sum_{l=1}^{2} u(\theta_l, q_l) - u(\theta_l, q_{l-1})$ since it could increase profits by charging this price, the resulting allocation still being implementable.

The same argument can be repeated for all the other feasible qualities. ■

3.6.3 Derivation of expression 3.6

Consider the profit function:

$$\Pi(\boldsymbol{\theta}) = \sum_{l=1,2,\ldots,N} \left[\sum_{m=1}^{l} \left(u(\theta_m, q_m) - u(\theta_m, q_{m-1}) \right) - c_l \right] \left[F(\theta_{l+1}) - F(\theta_l) \right]$$

Notice that this can be re-written as

$$\Pi(\boldsymbol{\theta}) = \sum_{l=1,2,\ldots,N} [u(\theta_l, q_l) - u(\theta_l, q_{l-1})][1 - F(\theta_l)] - c_l[F(\theta_{l+1}) - F(\theta_l)]$$

Consider the partial derivative with respect to θ_i and let $c_0 = 0$ to obtain

$$-[u(\theta_i, q_i) - u(\theta_i, q_{i-1})] f(\theta_i) + [u_\theta(\theta_i, q_i) - u_\theta(\theta_i, q_{i-1})][1 - F(\theta_l)] + (c_i - c_{i-1}) f(\theta_i)$$

This can be rearranged to obtain expression 3.6:

$$-[u(\theta_i, q_i) - u(\theta_i, q_{i-1}) - (c_i - c_{i-1})] f(\theta_i) + [u_\theta(\theta_i, q_i) - u_\theta(\theta_i, q_{i-1})][1 - F(\theta_l)]$$

3.6.4 Proof of proposition 3.6

The proof of proposition 3.6 is rather involved and derives from the combination of a number of individual results.

Lemma 3.9 *(i) Selling quality q_i is optimal if and only if $\lambda^*_{ij} = 0$ for all $j > i$. Selling quality q_i is not optimal if and only if $\lambda^*_{ij} < 0$ for at least a $j > i$. (ii) If $\lambda^*_{ij} < 0$, then $\lambda^*_{mj} < 0$ for every $m : i < m < j$.*

Proof. (i) Quality q_i is optimally sold if and only if $\theta^*_j > \theta^*_i$ for every $j > i$. Complementary slackness requires that $\lambda^*_{ij} = 0$ for every $j > i$. Similarly, quality q_i is not optimally sold if $\theta^*_j = \theta^*_i$ for at least a $j > i$. Complementary slackness requires that $\lambda^*_{ij} < 0$ for at least a $j > i$.

(ii) If $\lambda^*_{ij} < 0$ then $\theta^*_i = \theta^*_j$. Suppose there exists a quality $q_m : q_i < q_m < q_j$ such that $\lambda^*_{mj} = 0$. This implies that $\theta^*_j > \theta^*_m$. But, since $\theta^*_i = \theta^*_j$, this would also imply that $\theta^*_i > \theta^*_m$ which violates the monotonicity constraint that must hold in any implementable allocation. ■

Lemma 3.10 *(i) If quality q_i is optimally sold, then* $\sum_{l=1}^{i} \frac{\partial \Pi(\theta^*)}{\partial \theta_l} = 0$. *(ii) If quality q_i is not optimally sold, then* $\sum_{l=1}^{i} \frac{\partial \Pi(\theta^*)}{\partial \theta_l} > 0$

Proof. Consider the sum of the set of first-order-conditions for all qualities less or equal than q_i and notice that at the optimum

$$\sum_{l=1}^{i} \frac{\partial \Pi(\theta^*)}{\partial \theta_l} = -\sum_{j=1}^{i} \sum_{l=i+1}^{N} \lambda_{jl}^*$$

(i) If quality q_i is optimally sold, then lemma 3.9 states that $\lambda_{jl}^* = 0$ for all $l > j$. Therefore $\sum_{l=1}^{i} \frac{\partial \Pi}{\partial \theta_l} = 0$. (ii) If quality q_i is not optimally sold, then lemma 3.9 states that it must be the case that $\lambda_{jl}^* < 0$ for at least one $l : l > j$. Therefore, in this case $\sum_{l=1}^{i} \frac{\partial \Pi}{\partial \theta_l} > 0$. ■

Lemma 3.11 $\frac{\partial}{\partial \theta} \frac{\Delta S(\theta, q_j, q_i)}{\Delta S(\theta, q_i, q_{i-})} \leq (\geq) 0$ *if and only if* $\frac{\Delta S(\theta, q_j, q_i)}{\Delta S_\theta(\theta, q_j, q_i)} \geq (\leq) \frac{\Delta S(\theta, q_i, q_{i-})}{\Delta S_\theta(\theta, q_i, q_{i-})}$.

Proof. Consider the derivative

$$\frac{\partial}{\partial \theta} \left(\frac{\Delta S(\theta, q_j, q_i)}{\Delta S(\theta, q_i, q_{i-})} \right) \leq 0$$

$$\frac{\Delta S_\theta(\theta, q_j, q_i) \Delta S(\theta, q_i, q_{i-}) - \Delta S_\theta(\theta, q_i, q_{i-}) \Delta S(\theta, q_j, q_i)}{[\Delta S(\theta, q_i, q_{i-})]^2} \leq 0$$

i.e.

$$\Delta S_\theta(\theta, q_j, q_i) \Delta S(\theta, q_i, q_{i-}) - \Delta S_\theta(\theta, q_i, q_{i-}) \Delta S(\theta, q_j, q_i) \leq 0$$

i.e.

$$\frac{\Delta S(\theta, q_j, q_i)}{\Delta S_\theta(\theta, q_j, q_i)} \geq \frac{\Delta S(\theta, q_i, q_{i-})}{\Delta S_\theta(\theta, q_i, q_{i-})} \tag{3.10}$$

since by assumption $\Delta S_\theta(\theta, q_i, q_{i-}) > 0$. ■

Lemma 3.12 (Sufficient condition for creative destruction) *If, for every consumer type θ,*

$$\frac{\Delta S(\theta, q_j, q_i)}{\Delta S_\theta(\theta, q_j, q_i)} \geq \frac{\Delta S(\theta, q_i, q_{i-})}{\Delta S_\theta(\theta, q_i, q_{i-})}$$

then $\theta_j^* = \theta_i^*$.

Proof. I will show that quality q_i can be optimally sold alongside quality q_j, i.e. $\theta_j^* > \theta_i^*$, only if

$$\frac{\Delta S(\theta_i^*, q_j, q_i)}{\Delta S_\theta(\theta_i^*, q_j, q_i)} < \frac{\Delta S(\theta_i^*, q_i, q_{i-})}{\Delta S_\theta(\theta_i^*, q_i, q_{i-})}$$

Therefore, if for every θ

$$\frac{\Delta S(\theta, q_j, q_i)}{\Delta S_\theta(\theta, q_j, q_i)} \geq \frac{\Delta S(\theta, q_i, q_{i-})}{\Delta S_\theta(\theta, q_i, q_{i-})}$$

then it thus must be the case that quality q_i cannot be optimally sold alongside quality q_j, i.e. $\theta_j^* = \theta_i^*$.

If $\theta_j^* > \theta_i^*$, quality q_i can be optimally sold alongside quality q_j. If both quality q_j and quality q_i are optimally sold, then lemma 3.10 states that the following conditions must be satisfied

$$\sum_{l=1}^{j} \frac{\partial \Pi(\boldsymbol{\theta}^*)}{\partial \theta_l} = 0 \tag{3.11}$$

and

$$\sum_{l=1}^{i} \frac{\partial \Pi(\boldsymbol{\theta}^*)}{\partial \theta_l} = 0 \tag{3.12}$$

By substracting expression 3.12 from expression 3.11 one obtains

$$\sum_{l=i+1}^{j} \frac{\partial \Pi(\boldsymbol{\theta}^*)}{\partial \theta_l} = 0 \tag{3.13}$$

Let $\boldsymbol{\theta}_i^*$ the vector of marginal consumers such that $\theta_j = \theta_i^*$ for every $j > i$. Notice that

$$\sum_{l=i+1}^{j} \frac{\partial \Pi\left(\boldsymbol{\theta}_i^*\right)}{\partial \theta_l} = -\Delta S\left(\theta_i^*, q_j, q_i\right) f\left(\theta_i\right) + \Delta S_\theta\left(\theta_i^*, q_j, q_i\right)\left[1 - F\left(\theta_i^*\right)\right]$$

If $\frac{\partial^2 \Pi}{\partial \theta_i^2} < 0$ for every i, so that Kuhn and Tucker's necessary conditions for a maximum are also sufficient, it must be the case that

$$\sum_{l=i+1}^{j} \frac{\partial \Pi\left(\boldsymbol{\theta}_i^*\right)}{\partial \theta_l} > \sum_{l=i+1}^{j} \frac{\partial \Pi\left(\boldsymbol{\theta}^*\right)}{\partial \theta_l} \tag{3.14}$$

By combining expressions 3.13 and 3.14 one obtains that $\sum_{l=i+1}^{j} \frac{\partial \Pi(\boldsymbol{\theta}_i^*)}{\partial \theta_l} > 0$, i.e.

$$-\Delta S\left(\theta_i^*, q_j, q_i\right) f\left(\theta_i^*\right) + \Delta S_\theta\left(\theta_i^*, q_j, q_i\right)\left[1 - F\left(\theta_i^*\right)\right] > 0$$

i.e.

$$\frac{1 - F\left(\theta_i^*\right)}{f\left(\theta_i^*\right)} > \frac{\Delta S\left(\theta_i^*, q_j, q_i\right)}{\Delta S_\theta\left(\theta_i^*, q_j, q_i\right)} \tag{3.15}$$

Consider now quality q_{i-}, which is defined as the highest quality below q_i optimally sold by the monopolist. If quality q_{i-} is optimally sold, lemma 3.10 states that the following condition must be satisfied

$$\sum_{l=1}^{i-} \frac{\partial \Pi\left(\boldsymbol{\theta}^*\right)}{\partial \theta_l} = 0$$

and therefore, given condition 3.12, it must be the case that

$$\sum_{l=(i-)+1}^{i} \frac{\partial \Pi\left(\boldsymbol{\theta}^*\right)}{\partial \theta_l} = 0$$

Since no quality in between q_i and q_{i-} is optimally sold, and therefore $\theta_m = \theta_i^*$ for any $i > m \geq (i-)$ it follows that that

$$\sum_{l=(i-)+1}^{i} \frac{\partial \Pi(\boldsymbol{\theta}^*)}{\partial \theta_l} = -\Delta S(\theta_i^*, q_i, q_{i-}) f(\theta_i^*) + \Delta S_\theta(\theta_i^*, q_i, q_{i-}) [1 - F(\theta_i^*)]$$

This expression if equal to zero iff

$$\frac{1 - F(\theta_i^*)}{f(\theta_i^*)} = \frac{\Delta S(\theta_i^*, q_i, q_{i-})}{\Delta S_\theta(\theta_i^*, q_i, q_{i-})} \tag{3.16}$$

By combining expressions 3.15 and 3.16 it follows that a necessary condition for $\theta_j^* > \theta_i^*$ is that

$$\frac{\Delta S(\theta_i^*, q_j, q_i)}{\Delta S_\theta(\theta_i^*, q_j, q_i)} < \frac{\Delta S(\theta_i^*, q_i, q_{i-})}{\Delta S_\theta(\theta_i^*, q_i, q_{i-})}$$

Therefore, if for every θ

$$\frac{\Delta S(\theta, q_j, q_i)}{\Delta S_\theta(\theta, q_j, q_i)} \geq \frac{\Delta S(\theta, q_i, q_{i-})}{\Delta S_\theta(\theta, q_i, q_{i-})}$$

then $\theta_j^* = \theta_i^*$, i.e. selling quality q_i alongside quality q_j cannot not optimal. ■

Lemma 3.13 (Necessary condition for creative destruction) *If $\theta_j^* = \theta_i^*$ then*

$$\frac{\Delta S(\theta, q_j, q_i)}{\Delta S_\theta(\theta, q_j, q_i)} \geq \frac{\Delta S(\theta, q_i, q_{i-})}{\Delta S_\theta(\theta, q_i, q_{i-})}$$

Proof. I will show that, if for every consumer type θ

$$\frac{\Delta S(\theta, q_j, q_i)}{\Delta S_\theta(\theta, q_j, q_i)} < \frac{\Delta S(\theta, q_i, q_{i-})}{\Delta S_\theta(\theta, q_i, q_{i-})}$$

then $\theta_j^* > \theta_i^*$.

Suppose that $\theta_j^* = \theta_i^*$. In this case, lemma 3.10 states that the following neccessary conditions must be satisfied:

$$\sum_{l=1}^{j} \frac{\partial \Pi(\boldsymbol{\theta}^*)}{\partial \theta_l} = 0$$

and

$$\sum_{l=1}^{i} \frac{\partial \Pi(\boldsymbol{\theta}^*)}{\partial \theta_l} > 0$$

These two expressions imply that

$$\sum_{l=i+1}^{j} \frac{\partial \Pi(\boldsymbol{\theta}^*)}{\partial \theta_l} < 0$$

According to proposition 3.5, if $\theta_j^* = \theta_i^*$, then $\theta_j^* = \theta_i^* = \theta_m^*$ for every $m : j > m > i$ and therefore

$$\sum_{l=i+1}^{j} \frac{\partial \Pi(\boldsymbol{\theta}^*)}{\partial \theta_l} = -\Delta S(\theta_i^*, q_j, q_i) f(\theta_i^*) + \Delta S_\theta(\theta_i^*, q_j, q_i)[1 - F(\theta_i^*)]$$

For this expression to be negative it must be the case that

$$\frac{1 - F(\theta_i^*)}{f(\theta_i^*)} < \frac{\Delta S(\theta_i^*, q_j, q_i)}{\Delta S_\theta(\theta_i^*, q_j, q_i)} \tag{3.17}$$

Consider now q_{i-}, which is the highest quality below q_i sold in the market. If quality q_{i-} is optimally sold, then it is necessary that $\sum_{l=i}^{i-} \frac{\partial \Pi(\boldsymbol{\theta}^*)}{\partial \theta_l} = 0$. Therefore it must be the case that

$$\sum_{l=(i-)+1}^{i} \frac{\partial \Pi(\boldsymbol{\theta}^*)}{\partial \theta_l} > 0$$

because

$$\sum_{l=(i-)+1}^{i} \frac{\partial \Pi\left(\boldsymbol{\theta}^*\right)}{\partial \theta_l} = \sum_{l=1}^{i} \frac{\partial \Pi\left(\boldsymbol{\theta}^*\right)}{\partial \theta_l} - \sum_{l=1}^{i-} \frac{\partial \Pi\left(\boldsymbol{\theta}^*\right)}{\partial \theta_l}$$

where $\sum\limits_{l=1}^{i-} \frac{\partial \Pi(\boldsymbol{\theta}^*)}{\partial \theta_l} = 0$ and $\sum\limits_{l=1}^{i} \frac{\partial \Pi(\boldsymbol{\theta}^*)}{\partial \theta_l} > 0$.

Since no quality between q_i and q_{i-} is sold by definition of q_{i-}, $\theta_m^* = \theta_i^*$ for every $i > m > (i-)$ and therefore

$$\sum_{l=(i-)+1}^{i} \frac{\partial \Pi\left(\boldsymbol{\theta}^*\right)}{\partial \theta_l} = -\Delta S\left(\theta_i^*, q_i, q_{i-}\right) f\left(\theta_i^*\right) + \Delta S_\theta\left(\theta_i^*, q_i, q_{i-}\right)\left[1 - F\left(\theta_i^*\right)\right]$$

The condition $\sum\limits_{l=(i-)+1}^{i} \frac{\partial \Pi(\boldsymbol{\theta}^*)}{\partial \theta_l} > 0$ thus implies that

$$\frac{1 - F\left(\theta_i^*\right)}{f\left(\theta_i^*\right)} > \frac{\Delta S\left(\theta_i^*, q_i, q_{i-}\right)}{\Delta S_\theta\left(\theta_i^*, q_i, q_{i-}\right)} \tag{3.18}$$

Expression 3.17 and expression 3.18 together imply that a necessary condition for $\theta_j^* = \theta_i^*$ is that

$$\frac{\Delta S\left(\theta_i^*, q_j, q_i\right)}{\Delta S_\theta\left(\theta_i^*, q_j, q_i\right)} > \frac{\Delta S\left(\theta_i^*, q_i, q_{i-}\right)}{\Delta S_\theta\left(\theta_i^*, q_i, q_{i-}\right)}$$

Therefore, if for every θ

$$\frac{\Delta S\left(\theta, q_j, q_i\right)}{\Delta S_\theta\left(\theta, q_j, q_i\right)} < \frac{\Delta S\left(\theta, q_i, q_{i-}\right)}{\Delta S_\theta\left(\theta, q_i, q_{i-}\right)}$$

then it must be the case $\theta_j^* > \theta_i^*$.

This implies that $\theta_j^* = \theta_i^*$ only if

$$\frac{\Delta S\left(\theta, q_j, q_i\right)}{\Delta S_\theta\left(\theta, q_j, q_i\right)} \geq \frac{\Delta S\left(\theta, q_i, q_{i-}\right)}{\Delta S_\theta\left(\theta, q_i, q_{i-}\right)}$$

■

Proposition 3 follows from the combination of lemma 3.11, lemma 3.12 and lemma 3.13.

Lemma 3.12 states that if

$$\frac{\Delta S\left(\theta, q_j, q_i\right)}{\Delta S_\theta\left(\theta, q_j, q_i\right)} \geq \frac{\Delta S\left(\theta, q_i, q_{i-}\right)}{\Delta S_\theta\left(\theta, q_i, q_{i-}\right)}$$

Lemma 3.11 states that $\frac{\Delta S(\theta,q_j,q_i)}{\Delta S_\theta(\theta,q_j,q_i)} \geq \frac{\Delta S(\theta,q_i,q_{i-})}{\Delta S_\theta(\theta,q_i,q_{i-})}$ if and only if $\frac{\Delta S(\theta,q_j,q_i)}{\Delta S(\theta,q_i,q_{i-})}$ then $\theta_j^* = \theta_i^*$.

Lemma 3.11 states that $\frac{\Delta S(\theta,q_j,q_i)}{\Delta S_\theta(\theta,q_j,q_i)} \geq \frac{\Delta S(\theta,q_i,q_{i-})}{\Delta S_\theta(\theta,q_i,q_{i-})}$ if and only if $\frac{\Delta S(\theta,q_j,q_i)}{\Delta S(\theta,q_i,q_{i-})}$ is (weakly) decreasing in θ. Therefore, if $\frac{\Delta S(\theta,q_j,q_i)}{\Delta S(\theta,q_i,q_{i-})}$ is (weakly) decreasing in θ then $\theta_j^* = \theta_i^*$.

Similarly, lemma 3.13 states that if $\theta_j^* = \theta_i^*$ then

$$\frac{\Delta S\left(\theta, q_j, q_i\right)}{\Delta S_\theta\left(\theta, q_j, q_i\right)} \geq \frac{\Delta S\left(\theta, q_i, q_{i-}\right)}{\Delta S_\theta\left(\theta, q_i, q_{i-}\right)}$$

Combining this result with lemma 3.11 it follows that if $\theta_j^* = \theta_i^*$ then $\frac{\Delta S(\theta,q_j,q_i)}{\Delta S(\theta,q_i,q_{i-})}$ is (weakly) decreasing in θ.

Chapter 4

Optimal product line, secondary trading and used goods

4.1 Introduction

Second-hand trade is an important feature of many markets for durable goods. The organisation of secondary trading varies considerably, from relatively unstructured peer-to-peer transactions to developed markets in which a range of intermediaries purchase and resell used goods, match buyers and sellers, and often provide associated services such as quality certification and warranties.

A striking feature of many second-hand markets is that manufacturers of durable goods often act as intermediaries in the exchange of used goods. For instance, IBM was selling used IBM typewriters in the 60s and since 1987 it has been reselling used IT equipment, such as mainframes, midrange systems and PC servers.[1] Other manufacturers in IT industries, such as Compaq, Dell, Hewlett Packard, Cisco and SUN have recently started to offer trade-in promotions and to resell used equipment, which they usually refurbish and often adapt to the specific needs of customers. Often, the products that are resold by these manufacturers are acquired from patrons who traded in their used equipment in order to purchase a new one.[2] In the market for aircraft, Boeing has been selling even used Airbuses. Most car manufacturers sell certified pre-owned vehicles to which they extend warranty coverage, which now account for one third of used cars that are exchanged in the US.

In this paper I discuss a possible explanation of this commonly observed strategy of durable goods manufacturers. The main idea I put forward is that, by directly acting as an intermediary in the market for used good and by internalising the transaction costs of secondary trading, the manufacturer can induce former customers to upgrade while selling the new good to new consumers at the highest possible price.[3]

The model that I discuss in this Chapter describes a market for a perfectly durable good subject to quality-improving technological change, which is analysed over a two-period horizon. Each period, a new cohort of heterogeneous consumers enters the market. Consumers in each cohort differ in their willingness to pay for quality. In addition, in period 2 an additional source of heterogeneity arises because consumers with the same willingness to pay

[1] Howevcer, notice that in the market for personal computers that I studied in Chapter 1 and in Chapter 2 sales of used goods by PC manufacturers is not a common phenomenon.

[2] In some cases, when a product under warranty is in need of service, manufacturers prefer to replace the product with a new one altogether. The products acquired are often refurbished and (re)sold at discounted prices. This phenomenon is different and complementary to the one that I investigate in this Chapter, which is related to consumers upgrading to new better durable goods and not to consumers requesting assistance for faulty products.

[3] Notice that I use the term "transaction costs" to indicate the costs that must be borne in order for a used good to be exchanged. This usage is consistent with the literature on secondary trading that I review in the next section.

for quality may have different purchase histories: some consumers have just entered the market while other consumers may hold a durable good purchased in period 1.

In period 2, the manufacturer may find it profitable to sell the new technology to consumers in both cohorts, i.e. to consumers with different purchase histories, but its ability of doing so depends on the nature of the good. For some durable goods such as software, which I term "knowledge-based", the manufacturer can profitably sell an "upgrade" to former consumers because consumers value the upgrade more than its cost. In this case, the manufacturer can optimally offer a trade-in promotion to discriminate between consumers with a different purchase history and charge former consumers a price that reflects only the incremental utility that they would enjoy from purchasing the new technology while charging new consumers the full price.[4]

In other industries, which I term "industrial", however, the manufacturer cannot profitably sell the new good to old consumers if a secondary market does not exist because former consumers value the upgrade to the new technology less than its cost.

The fact that used goods are traded has two effects on the market for the new technology. Used goods are a substitute for new goods and hence tend to decrease the price at which the new technology can be sold. At the same time, however, second-hand trade allows former patrons to upgrade to the new technology or, in general, to upgrade more frequently.

Secondary trading often entails possibly significant transaction costs deriving, for instance, from the effort required to match buyers and sellers or from the provision of quality certification to address problems of asymmetric

[4]Buy-back strategies have also been studied by Levinthal and Purhoit (1989), Lee and Lee (1998) and Fudenberg and Tirole (1998).

information.

The existence of transaction costs implies that the price at which the new good can be sold to former patrons is lower than the price at which the new good can be sold to new consumers. In fact, former patrons choose between buying the new good and keeping the one in their possession, which entails an implicit cost equal to the price at which the used good can be sold in the secondary market. However, new consumers choose instead between buying the new good and buying a used good at the market price which is higher than the price at which the used good can be sold by former consumers. If there is a competitive market for used goods, in fact, the difference between the price at which a used good can be sold and the price at which it can be purchased by consumers is equal to the transaction cost that are borne when trading a used good. Transaction costs therefore drive a wedge between the price that new and former consumers are willing to pay for the new technology. As a result, if the manufacturer wants to sell the new technology to both groups of consumers, it is constrained by the lower price that former patrons would be willing to pay for the new technology, given the price that they can obtain for their used good.

This suggests that there is scope for the manufacturer to improve profits by offering a trade-in promotion and internalising transaction costs. A trade-in promotion allows the manufacturer to sell the new technology to both former patrons and to new consumers at the highest price possible. This is because, rather than allowing former consumers to upgrade by lowering the price at which the new good is sold, the monopolist allows them to upgrade by purchasing their used goods at a higher price than the one offered by competitive traders. This strategy does not affect the profitability of a transaction with former consumers who upgrade to the new technology: what the monopolist gives with one hand it takes with the other. However, charging a high price for the new good increases the profits from sales to new consumers.

This is not to say that the existence of a secondary market always increases the profits of the monopolist. In fact, the monopolist may be better off by simply selling the new durable good to new consumers at a high price. However, when a secondary market exists, the monopolist is better off by directly trading used goods and internalising transaction costs..

In the remainder of this Chapter I first discuss the related literature. In section 4.3, I present the model and the analysis of the equilibrium. Final conclusions follow.

4.2 Related literature

Two main strands of the literature on durable goods are closely related to this study. The first addresses the long-standing question of what impact the existence of an active second-hand market has on the profitability of a monopolist seller in the primary market. The second focuses specifically on pricing and product-line decisions of a monopolist who sells successive technological generations of a durable good and considers the optimality of trade-ins and buy-backs.

The impact that the existence of a second-hand market has on a monopolist seller has been one of the major questions posed in the economic literature on durable goods. Swan (1970, 1972, 1980) suggested that the existence of a second-hand market does not constrain the profits of the monopolist in the primary market: the seller can charge a price that reflects the net present value of the stream of services provided by the durable good to possibly a number of future owners. This striking result hinges on a number of restrictive assumptions, and most notably consumers homogeneity and perfect susbstitutability of the service flow provided by durable goods of different vintages. It is not surprising that in richer settings, Swan's result does not hold and that second-hand trade affects the profitability of the monopolist seller in the primary market.

There are various effects at stake, however, and their balance is not clear. On the one hand, used goods are a substitute for new goods, which tends to decrease the profitability of the monopolist seller. On the other hand, consumers may upgrade (more often) to new goods if they are able to sell their used goods and would be willing to pay more for goods that would command a higher resale price. These ideas have been rigorously developed in models where consumers are heterogeneous and durable goods are subject to physical quality deterioration or technological progress.

Anderson and Ginsburgh (1994) consider a model with heterogeneous consumers, quality depreciation and positive transaction costs in the second-hand market. They suggest that existence of a smooth second-hand market may increase the price at which the new good can be sold since consumers are able to anticipate a higher resale value and the monopolist may benefit from it because it achieves an indirect form of indirect price discrimination between consumers with different willingness to pay for quality. Anderson and Ginsburgh show that, in any equilibrium where there is an active secondary market with a positive price for the used good, the monopolist's profits are decreasing in the level of transaction costs. Nevertheless, the manufacturer may be better off if the second-hand market does not exist. This is a global optimum for the manufacturer when individuals place a high value on newness because the monopolist would benefit little from the increase in demand that arises due to consumers benefiting from resale value. Like Anderson and Ginsburgh (1994), I also consider a model with transaction costs in the second-hand market.

However, I do not focus on the effect that transaction costs have on profits because of consumers being able to anticipate a higher resale value, but on the effect that transaction costs have on the monopolists' profits from selling the new generation of the durable good to consumers with different purchase histories, and I explain why it may be profitable for the monopolist to actually internalise transaction costs. Unlike Anderson and Ginsburgh, in fact, I study the benefits and the costs of reducing transaction costs in the secondary market.

Hendel and Lizzeri (1999) consider a model similar to Anderson and Ginsburgh's and focus on the ability that a manufacturer has to affect the product's durability. They consider only the case of frictionless exchange of used goods and argue that the existence of a second-hand market would never hurt the manufacturer. The explanation of this result relies on the notion that the monopolist would be able to achieve a form of indirect price discrimination through the second-hand market. Hendel and Lizzeri generalise the paper by Waldman (1996) who argues that the monopolist may want to eliminate the existence of the second-hand market by using leasing. Waldman had considered the special case where there are only two types of consumers and the monopolist could not profitably sell to those consumers with lower willingness to pay for quality.

Porter and Sattler (1999) consider a similar model with heterogeneous consumers and product depreciation, which is applied to the study of patterns of trade in second-hand markets. Unlike Hendel and Lizzeri (1999) they consider also the case in which there are transaction costs in the second-hand market. Their main conclusion is that durable goods that depreciates relatively quickly have steeper price declines, but higher volumes of trade.

Kumar (2002) considers a model in which a durable good monopolist strategically chooses both the price and the quality of a durable good over time and shows that the monopolist exploits resale trading by consumers to intertemporally price discriminate by strategically controlling the rate of technological obsolescence.

The main difference between this paper and these studies is that I consider second-hand markets in a setting where there is technological change rather than quality depreciation and that I explicitly consider the possibility that the manufacturer discriminates between consumers with different purchase histories. Moreover, this paper considers subperfect game equilibria and not only the case where the monopolist commits to a path of prices/output.

These features make this study close to another strand of the literature which considers the issue of pricing and product-line decisions in durable goods markets subject to technological progress. These models discuss the relationship between static quality discrimination and intertemporal discrimination and the profitability of trade-in and buy-back strategies. Unlike this study, however, they do not emphasise the role of second-hand markets in making upgrading possible and normally consider the case of a frictionless secondary market (or no secondary market at all).

Levinthal and Purhoit (1989) consider a two-period model of a durable-good monopolist who can produce a good of higher quality in the second period. They consider linear demand, no production costs and a frictionless second-hand market. They discuss the profitability of a buy-back strategy by which the monopolist kills off the second-hand market.

Lee and Lee (1998) consider a two-period model of endogenous technological change in a durable good market where there are discrete consumers' types and no second-hand market (i.e. consumer scrap the used good in their possession to buy new). They study the profitability of an upgrade policy when the monopolist discriminates both between consumers with different purchase histories and between consumers of different types.

Fudenberg and Tirole (1998) consider the relationship between intertemporal price discrimination and static quality discrimination in a model similar to Lee and Lee's. Unlike Lee and Lee, however, they consider exogenous technological change, a continuous distribution of consumer types, and both the case of a frictionless second-hand market and of no secondary trading.

I study a model with a structure similar to Lee and Lee's and Fudenberg and Tirole's but, like Anderson and Ginsburgh (1994), I consider transaction costs and I propose a new explanation for the role of trade-in promotions.

4.3 The model

Time is discrete and there are only two periods, $t = 1, 2$.[5] Each period, a new cohort of consumers of mass $2n_t$ enters the market. In each cohort there are two groups of consumers with different willingness to pay for quality θ_H, θ_L, with $\theta_H > \theta_L$, each of mass n_t. Each consumer purchases at most one unit of the good at any date and a consumer of type θ who consumes a good of quality V obtains a per-period utility equal to $u(\theta, V) = \theta V$.

In period 1, the monopolist produces a perfectly durable good of quality V_1 at constant marginal cost c_1. In period 2, the monopolist can sell a new generation of the good of quality V_2, with $V_2 > V_1$, which is produced at marginal cost c_2. I assume that the cost of "upgrading" a good of quality V_1 to quality V_2 is the same as the cost of producing the good from scratch, i.e. c_2. I also assume that $\theta_L V_t < c_t$, $t = 1, 2$, i.e. a new durable good can never be profitably sold to θ_L consumers and that $\theta_H V_2 - c_2 > \theta_H V_1 - c_1$ so that

[5]This two-period setup is similar to the one used by Fudenberg and Tirole (1998).

selling technology V_2 is more profitable than selling technology V_1.

If θ_H consumers purchase good V_1 in period 1, there is an opportunity for second-hand trade in period 2, since the used good can be potentially sold to θ_L consumers.

I assume that used goods cannot be exchanged by consumers without the intervention of a trader, who incurs the transaction cost τ for each unit exchanged.[6] The market for trading used goods is perfectly competitive and hence traders are price-takers. The monopolist can also intervene in the second-hand market by purchasing used goods and reselling them, or by acquiring used goods and disposing of them. I denote by p_{bu} the price at which used goods are bought by competitive traders and by p_{su} the price at which used goods are sold to consumers. I denote by p_{bu}^* the price at which used goods are purchased by the monopolist. This price can potentially be different from p_{bu} because the monopolist is not subject to the zero-profit condition that must hold for the perfectly competitive traders.

In the analysis, I distinguish between two categories of goods: "industrial" goods and "knowledge-based" goods. The definition is based on the comparison between the value of the upgrade $(V_2 - V_1)$ to θ_H consumer and the marginal cost of the new good (which is equal to the marginal cost of the upgrade). An industrial good is defined as a good for which $\theta_H (V_2 - V_1) < c_2$, i.e. θ_H consumers value the upgrade less than its marginal cost. This implies that in any equilibrium in which consumers who bought technology V_1 upgrade to technology V_2 a secondary market must be active because the manufacturer cannot profitably sell them an upgrade if the good is not resold to

[6] Notice that the model is equivalent to one in which used goods are exchanged directly by consumers and in which the buyer must pay cost τ on top of the amount given to the seller. The key assumption is that there is a wedge between the price received by a seller and the price paid by a buyer when a used good is exchanged.

θ_L consumers. This feature may depict well markets for important durable goods such as cars and some expensive IT equipment.

A knowledge-based good is instead defined as one for which $\theta_H (V_2 - V_1) \geq c_2$, i.e. θ_H consumers value the upgrade more than the marginal cost of the new good, i.e. the marginal cost of the upgrade. In this case, upgrading can occur even if there is no active secondary market and I will show that, in fact, the monopolist may have the incentive to kill-off secondary trading. Software and text-books are good examples of knowledge-based goods.

4.4 Analysis

Since I am interested in studying the intervention of the manufacturer in the secondary market, I will focus the analysis on period 2, assuming that in the first period the monopolist had sold technology V_1 to n_1 consumers of type θ_H and therefore that a secondary market can indeed be potentially active in period 2.

Since, by assumption, selling technology V_2 is more profitable than selling technology V_1 and since since the monopolist cannot profitably sell (new) technology V_1 to θ_L consumers, the monopolist does not optimally produce technology V_1 in period 2. In period 2, therefore, the manufacturer sells only new durable goods of quality V_2, either to new θ_H consumers in cohort 2 only or to θ_H consumers in both cohort 1 and cohort 2.

I first analyse the case in which the monopolist sells to both cohorts of consumers, i.e. the equilibrium with upgrading, and then the case in which the monopolist sells V_2 to consumers in cohort 2 only, i.e. the equilibrium with no upgrading. I determine optimum profits within each regime and then compare the two regimes to determine the global optimum for the monopolist.

4.4.1 Equilibrium with upgrading

In this section I consider the case in which the new durable good V_2 is sold to consumers in both cohorts. I first consider the case in which the monopolist sells the new durable good of quality V_2 to consumers in both cohorts and the second-hand market is active. An active second-hand market is a necessary condition for the new technology V_2 to be profitably sold to θ_H consumers in cohort 1 in the case of an industrial good. This is because a consumer who purchased V_1 in period 1 is willing to pay only $\theta_H (V_2 - V_1)$ to purchase good V_2, but in the case of an industrial good this is lower than the marginal cost of purchasing the "upgrade", i.e. $\theta_H (V_2 - V_1) < c_2$. I will then consider the case in which the monopolist sells a knowledge-based good to consumers in both cohorts even if there is not active secondary market.

In any equilibrium in which the new technology is purchased both by consumers in cohort 2 and by consumers in cohort 1 and there is an active secondary market two incentive constraints must be satisfied.

A θ_H consumer in cohort 2 would purchase technology V_2 only if the utility he would get by doing is at least equal to the utility he would obtain by purchasing a used good V_1 at price p_{su}:

$$\theta_H V_2 - p_2 \geq \theta_H V_1 - p_{su}$$

which can be rewritten as

$$p_2 \leq \theta_H (V_2 - V_1) + p_{su} \tag{4.1}$$

Consumers θ_H in cohort 1 would instead purchase good V_2 only if upgrading yields a higher utility than than keeping the old technology V_1:

$$\theta_H V_2 - p_2 + \max\{p_{bu}^*, p_{bu}\} \geq \theta_H V_1$$

where p_{bu} and p_{bu}^* are the prices at which a consumer in cohort 1 can sell his used good to a competitive trader and to the monopolist respectively. This constraint can be rewritten as

$$p_2 \leq \theta_H (V_2 - V_1) + \max\{p_{bu}^*, p_{bu}\} \tag{4.2}$$

If there are no transaction costs in the competitive second-hand market, the two constraints coincide since the price at which the used good is purchased by traders would be equal to the price at which the used good is resold to consumers. However, the existence of transaction costs drives a wedge between the price at which used goods are purchased by competitive traders p_{bu} and the price at which used goods are sold to consumers p_{su}. In turn, this drives a wedge between the highest price at which the new good can be sold to new consumers and the highest price at which the new good can be sold to former patrons. This can be seen by collapsing constraint 4.1 and constraint 4.2 into the following expression:

$$p_2 \leq \theta_H (V_2 - V_1) + \min\{p_{su}, \max\{p_{bu}^*, p_{bu}\}\}$$

which shows that the highest price at which the new good can be sold when there is an active secondary market is constrained by the minimum between the price at which used goods are purchased and the price at which they are sold. Notice that, for any p_{bu}^* and p_{bu}, it is optimal for the manufacturer to set p_2 as large as possible, i.e.

$$p_2 = \theta_H (V_2 - V_1) + \min\{p_{su}, \max\{p_{bu}^*, p_{bu}\}\}$$

The following proposition states that, in any equilibrium in which the secondary market is active, it is always profitable for the monopolist to intervene directly in the exchange of used goods and to internalise transaction costs.

Proposition 4.1 *In any equilibrium in which the monopolist sells the new technology to θ_H consumers in both cohorts and there is an active secondary market, it is optimal for the monopolist to intervene directly in the secondary market and to purchase used goods at price $p_{bu}^* = \theta_L V_1$, i.e. to internalise transaction costs.*

Proof. First of all, notice that, in any equilibrium in which θ_H consumers in both cohorts purchase the new durable good in period 2, because of the assumption of perfect competition, the highest price that a competitive trader can offer to purchase a used good is $p_{bu} = \theta_L V_1 - \tau$ and the highest price at which it would sell the good is $p_{bu} = \theta_L V_1$.

I will now consider two possible pricing strategies by the monopolist: (i) $p_{bu}^* \leq (\theta_L V_1 - \tau)$, (ii) $p_{bu}^* > \theta_L V_1 - \tau$. In strategy (i) the monopolist offers a price to purchase used goods which is equal or lower than the price offered by competitive traders. In strategy (ii) the monopolist purchases used goods at a price higher than the price offered by competitive traders. I will show that strategy (ii) is more profitable than strategy (i) and, more specifically, that it is optimal for the monopolist to set $p_{bu}^* = \theta_L V_1$, i.e. internalise transaction costs τ.

(i) If the monopolist offers price $p_{bu}^* \leq (\theta_L V_1 - \tau)$, the price that a consumer in cohort 1 would receive for her used good is $(\theta_L V_i - \tau)$. In this case, the highest price that the manufacturer can charge for the new good and sell it to consumers in both cohorts is

$$p_2 = \theta_H (V_2 - V_1) + (\theta_L V_1 - \tau)$$

because the incentive constraint for consumers in cohort 1 is binding. This results in period-2 profits equal to

$$\Pi_2 = (n_1 + n_2) [\theta_H (V_2 - V_1) + (\theta_L V_1 - \tau) - c_2] \tag{4.3}$$

(ii) If the monopolist sets price $p^*_{bu} = \widetilde{p}^*_{bu} > \theta_L V_1 - \tau$, competitive traders would not be active and the monopolist would manage all secondary trading. In this case the monopolist's period-2 profits would be

$$\widetilde{\Pi}_2 = (n_1 + n_2) [\theta_H (V_2 - V_1) + \min(\widetilde{p}^*_{bu}, \theta_L V_1) - c_2] + n_1 [\theta_L V_1 - \widetilde{p}^*_{bu} - \tau]$$

The monopolist would sell the new durable good to $(n_1 + n_2)\, \theta_H$ consumer at price $[\theta_H (V_2 - V_1) + \min(\widetilde{p}^*_{bu}, \theta_L V_1)]$ and it would trade n_1 used goods purchased at price $\widetilde{p}^*_{bu}$ and sold at price $\theta_L V_1$ incurring the cost τ for each used good traded.

Notice that is optimal for the monopolist to purchase used goods at price $\widetilde{p}^*_{bu} = \theta_L V_1$. Notice in fact, that for any $\widetilde{p}^*_{bu} > \theta_L V_1$, $\frac{\partial}{\partial \widetilde{p}^*_{bu}} \widetilde{\Pi}_2 = -n_1 < 0$ and that for any $\widetilde{p}^*_{bu} : \theta_L V_1 \geq \widetilde{p}^*_{bu} \geq \theta_L V_1 - \tau$, $\frac{\partial}{\partial \widetilde{p}^*_{bu}} \widetilde{\Pi}_2 = n_2 > 0$. It follows that, if the monopolist is to purchase used goods at a price higher than the price offered by competitive traders, it would optimally set $\widetilde{p}^*_{bu} = \theta_L V_1$, i.e. it would optimally internalise transaction costs. This allows the monopolist to obtain profits equal to

$$\widetilde{\Pi}_2 = (n_1 + n_2)\left[\theta_H (V_2 - V_1) + \theta_L V_1 - c_2\right] - n_1 \tau \tag{4.4}$$

By comparing now expression 4.3 and expression 4.4 it is straightforward to notice that $\widetilde{\Pi}_2 > \Pi_2$, which implies that the monopolist is always better off by intervening in the secondary market and internalising transaction costs rather than by leaving secondary trading to be managed by competitive traders. ■

The intuition for this result is the following. The monopolist can induce former patrons to upgrade by either selling the new good at a low price or by purchasing their used good at a high price. The profits obtained from former consumers who upgrade are the same under the two strategies: what the monopolist gives with one hand, it takes with the other. However, the two strategies are not equivalent in terms of profits from sales to new consumers: the monopolist obtains higher profits when charging the highest price possible for the new good to consumers in cohort 2. By offering trade-ins and internalising transaction costs, the manufacturer can maintain a high price for the new good while inducing former patrons to upgrade. Intervening directly in the secondary market is therefore a way to ease the transaction that allows a consumer who purchased the old technology to upgrade to the new one and to sell the new technology at the highest price possible to consumers who purchase the good for the first time.

In the case of an industrial good, the existence of a second-hand market is a necessary condition for consumers in cohort 1 to upgrade to technology V_2 in equilibrium. This is because, by assumption, the manufacturer of an industrial good cannot profitably sell an upgrade to former patrons who value the upgrade less than its cost of production. In the case of a knowledge-based good, however, the monopolist can profitably sell the new technology to θ_H

consumers in both cohorts even if there is no active secondary market. I will show that, in this case, the monopolist can actually be better off by purchasing the used good and scrapping it, thereby practically killing-off the secondary market. In this case the monopolist purchases the used good simply as a means to discriminate between consumers with a different purchase history: it can charge consumers in cohort 1 a low price that reflects only their willingness to pay to upgrade, while charging consumers in cohort 2 the full price. This is possible because, by killing-off the secondary market, the monopolist does not create any substitution opportunity for new consumers in cohort 2.

Proposition 4.2 *If (and only if) $\frac{\theta_H}{\theta_L} > \frac{n_1+n_2}{n_2}$, it is optimal for the manufacturer of a knowledge-based good to offer a trade-in promotion and to kill-off the secondary market .*

Proof. If the monopolist kills off the secondary market, all θ_H consumers in both cohorts purchase the new technology and the monopolist is able to discriminate between the two cohorts and to extract all consumer surplus by charging price $p_2 = \theta_H V_2$ to consumers who do not have a used good to trade-in and price $p_2 = \theta_H (V_2 - V_1)$ to consumers who can trade-in a used good. By doing so, and assuming that the manufacturer would still incur the transaction cost τ (which can be interpreted as a disposal cost), it would get profits equal to

$$\begin{aligned}\Pi_1 &= n_1 [\theta_H (V_2 - V_1) - c_2 - \tau] + n_2 [\theta_H V_2 - c_2] = \\ &= (n_1 + n_2) [\theta_H V_2 - c_2] - n_1 (\theta_H V_1 + \tau)\end{aligned}$$

If the monopolist trades used goods (which I showed that is always more profitable than allowing secondary trading by competitive traders), it obtains profits

$$\Pi_2 = (n_1 + n_2) [\theta_H (V_2 - V_1) + \theta_L V_1 - c_2] - n_1 \tau$$

The difference $\Pi_1 - \Pi_2$ is equal to

$$\Pi_1 - \Pi_2 = (n_1 + n_2) (\theta_H - \theta_L) V_1 - n_1 \theta_H V_1$$

which is positive if and only if $\frac{\theta_H}{\theta_L} > \frac{n_1 + n_2}{n_2}$. ■

In the case of a knowledge-based good, the monopolist can sell the new technology to consumers in both cohorts by either easing or killing-off secondary trading. Killing-off the secondary market is akin to paying $n_1 \theta_H V_1$ (by offering a discount to consumers in cohort 1) in order to capture the full surplus of high-type consumers from consuming V_2 rather than leaving them information rents. It is therefore optimal if consumers' willingness to pay for quality is relatively heterogeneous or if there are not too many former consumers.

4.4.2 Equilibrium with no upgrading

I have characterised the optimal strategy of the monopolist when it sells the new technology to both cohorts of consumers. It is also possible for the monopolist to sell the new durable good to consumers in cohort 2 only. In this case it would obtain a profit equal to $n_2 [\theta_H V_2 - c_2]$.

Proposition 4.3 identifies the condition(s) under which it is optimal for the monopolist to sell the new technology to both groups of consumers rather than to consumers in cohort 2 only. This is a substantial question relevant only for the manufacturer of an industrial good since, as long as transaction costs are not too large, in the case of a knowledge-based good the monopolist is always worse off by not offering a trade-in promotion and selling to consumers in cohort 2 only because it can perfectly price discriminate between consumers with different purchase histories by killing-off the secondary market.

Proposition 4.3 *(i) The manufacturer of an "industrial" good optimally sells the new technology to consumers in both cohorts if and only if V_2 is large enough. (ii) Provided that transaction/disposal costs are low enough, the manufacturer of a "knowledge-based" good would always optimally sell the new durable good to both cohorts of consumers.*

Proof. Consider first the case of an industrial good. The profits that the monopolist obtains when it sells the new technology V_2 to both groups of consumers and by intervening directly in the secondary market are

$$\widetilde{\Pi}_2 = (n_1 + n_2)\left[\theta_H (V_2 - V_1) + \theta_L V_1 - c_2\right] - n_1 \tau$$

The profits that the monopolist obtains when selling V_2 only to consumers in cohort 2 are given by

$$\widehat{\Pi}_2 = n_2 \left[\theta_H V_2 - c_2\right]$$

It is optimal to sell to both groups of consumers if and only if $\widetilde{\Pi}_2 \geq \widehat{\Pi}_2$, i.e. iff

$$(n_1 + n_2)\left[\theta_H (V_2 - V_1) + \theta_L V_1 - c_2\right] - n_1 \tau \geq n_2 \left[\theta_H V_2 - c_2\right]$$

This is true if and only if

$$V_2 \geq \frac{(\theta_H - \theta_L)(1 + n_2/n_1) V_1 + (c_2 + \tau)}{\theta_H}$$

Consider now the case of a knowledge-based good. By offering a trade-in promotion the manufacturer obtains profits equal to

$$\Pi_2 = n_1 \left[\theta_H (V_2 - V_1) - c_2 - \tau\right] + n_2 \left[\theta_H V_2 - c_2\right]$$

since it would sell the new durable good to n_1 consumers in cohort 1 at price $\theta_H (V_2 - V_1)$ bearing a cost equal to $(c_2 + \tau)$ for each unit sold and it would sell the new durable good to n_2 consumers in cohort 2 at price $\theta_H V_2$ bearing a cost of c_2 for each unit sold.

However, by selling only to θ_H consumers in cohort 2, the manufacturer obtains profits equal to

$$\widehat{\Pi}_2 = n_2 \left[\theta_H V_2 - c_2\right]$$

Offering a trade-in promotion is therefore optimal as long as

$$\left[\theta_H (V_2 - V_1) - c_2 - \tau\right] \geq 0$$

i.e. $\tau < \theta_H (V_2 - V_1) - c_2$. This is simply stating the a θ_H-type consumer must be willing to pay a price for the upgrade at least equal to the total cost of the upgrade, i.e. the marginal cost of the new technology plus the transaction/disposal cost τ. ■

The intuition for this result is the following. In the case of an industrial good, if the monopolist sells the new technology only to new consumers, it can extract all surplus of θ_H consumers in cohort 2 from consuming technology V_2. By selling the new technology to former patrons as well, the monopolist can capture some surplus that θ_H-type consumers in cohort 1 derive from the upgrade but it must leave information rents to θ_2 consumers in cohort 2. The benefit is increasing in the magnitude of technological progress, while the cost is not. Hence, when technological progress is sufficiently fast in the sense that

V_2 is sufficiently large, the manufacturer of an industrial good would prefer to sell the new technology to both new and former patrons and have an active secondary market.[7]

This result may explain why trade-in promotions are more common in innovative markets than in more static industries. For instance, trade-in promotions and reselling are very rare in markets for white goods, which are characterised by a rather slow rate of technological progress. By contrast, trade-in offers are now widespread in many dynamic IT industries, from personal computers to mainframes and routers in which manufacturers normally resell the used goods that they acquire from consumers who upgrade to new technologies.

4.5 Conclusions

Secondary trading plays an important role in many markets for durable goods since it allows consumers to upgrade, or to upgrade more often, to new durable goods of higher quality. There is evidence that in many industries manufacturers of durable goods often act as intermediaries in the trade of used goods, reselling used goods that are generally acquired from consumers who upgrade to the new technology.

The model presented in this paper offers an explanation of why this can be a profitable strategy. The manufacturer has the incentive to intervene directly in the secondary market and to internalise transaction costs in order to ease the process that allows former patrons to upgrade to the new technology and to be able to sell the new technology to both new and former consumer at the

[7]Notice that this result is different to the one obtained by Anderson and Ginsburgh (1994) who show that the monopolist prefers not to have an active secondary market when individuals place a high value on newness.

highest price possible.

Transaction costs in the second-hand market reduce the price at which the new good can be sold to former patrons compared to the price at which the new good could be sold to new consumers. If the monopolist sells the new durable good to both groups of consumers, it is constrained by the price that it can charge to former patrons. By internalising transaction costs the manufacturer can set a higher price for the new good while allowing former patrons to upgrade by purchasing their used goods at a higher price than the one offered by external traders. This is more profitable than inducing upgrading by lowering the price of the new good because it allows the monopolist to charge a higher price for the new durable good sold to new consumers in the market.

For some durable goods, the existence of a second-hand market is not necessary for former patrons to upgrade to the new better technology since the monopolist can profitably sell an upgrade to former consumers. In this case, the manufacturer may optimally buy-back used goods, i.e. it may implicitly purchase (and dispose of) used goods from former consumers by offering a discount to those consumers who trade in their used good. In this case the manufacturer can price discriminate between consumers with different purchase histories by killing-off the secondary market.

References

Abel A and J. Eberly (1994), "A Unified Model of Investment under Uncertainty", *American Economic Review*, 84, pp.1369-84.

Alexander, P. (1997), "Product variety and market structure: A new measure and a simple test", *Journal of Economic Behaviour and Organization*, 32, pp.207-15.

Anderson S. and V. Ginsburgh (1994), "Price discrimination via second-hand markets, *European Economic Review*, 38, pp.23-44.

Attanasio, O. (2000), "Consumer durables and inertial behaviour: estimation and aggregation of (S,s) rules for automobile purchases", *Review of Economic Studies*, 67(4), pp.667-96.

Bajari P., C. L. Benkard (2005), "Demand estimation with heterogeneous consumers and unobserved product characteristics: a hedonic approach", *Journal of Political Economy*, 113(6), pp.1239-76.

Bajari P., C. L. Benkard and J. Levin (2007), "Estimating dynamic models of imperfect competition", *Econometrica*, 5, pp.1331-70, Sep 2007.

Bayus, B. and W. P. Putsis (1999), "Product proliferation: An empirical analysis of product line determinants and market outcomes", *Marketing Science*, 18(2), pp.137-53.

Benkard CL (2004), "A dynamic analysis of the market for wide-bodied commercial aircraft", *Review of Economic Studies*, 71(3), 581-611.

Berry S. (1992), "Estimation of a model of entry in the airline industry", *Econometrica*, 60(4), pp.887-917.

Berry, S., Levinsohn, J., and A. Pakes (1995), "Automobile prices in market equilibrium", *Econometrica*, 63(4), pp.841-90.

Berry, S. and J. Waldfogel (2001), "Do mergers increase product variety? Evidence from radio broadcasting", *Quarterly Journal of Economics*, 116(3), pp.1009-25.

Berry, S. and A. Pakes (2007), "The pure characteristics demand model", *mimeo.*

Besanko D and U Doraszelski (2004), "Capacity dynamics and endogenous asymmetries in firm size", *Rand Journal of Economics*, 35(1), pp.23-49.

Brynjolfsson, E., Smith, M., and Y. Hu (2003), "Consumer surplus in the digital economy: estimating the value of increased product variety at online booksellers", *Management Science*, 49(11).

Bresnahan, T. and P. Reiss (1990), "Entry in Monopoly Markets", *Review of Economic Studies*, 57, pp.531-553.

Bresnahan, T. and P. Reiss (1991), "Empirical models of discrete games", *Journal of Econometrics*, 48, pp.57-82.

Bresnahan, T. and P. Reiss (1991), "Entry and Competition in Concentrated Markets," *Journal of Political Economy*, 99, pp.977-1009.

Bresnahan, T.F. Stern S. and M. Trajtenberg (1997) "Market Segmentation and the Sources of Rents from Innovation: Personal Computers in the Late 1980s", *RAND Journal of Economics*, 28, pp. S17-44.

Caballero RJ and EM Engel (1993), "Microeconomics adjustment hazards and aggregate dynamics", *Quarterly Journal of Economics*, 108(2), pp.359-83.

Caballero RJ and EM Engel (1999), "Explaining investment dynamics in US manufacturing: a generalisd (S,s) approach", *Econometrica*, 67(4), p.783-

826.

Caplin A and D Spulber (1987), "Menu costs of the neutrality of money", *Quarterly Journal of Economics*, 102, pp.703-25.

Carranza, J. (2006), "Product innovation in durable goods markets", *mimeo*, University of Wisconsin-Madison.

Chisholm and Norman (2005), "When to exit a product: evidence from the U.S. motion pictures exhibition market", Department of Economics Working Paper, Tufts university.

Crawford G. and M. Shum (2007), "Monopoly quality degradation and regulation in cable television", *Journal of Laws and Economics*, 50, pp.181 - 219.

Davis, P. (2006), "Quantity competition in the presence of indivisibilities and heterogeneous firms", *Journal of Econometrics*, 134(1), pp. 187-214.

Deltas, G. and E. Zacharias (2006) "Entry Order and Pricing over the Product Cycle: The Transition from the 486 to the Pentium Processor", *International Journal of Industrial Organization*, 24(5),pp. 1041-69.

Dixit A. and J. Stiglitz (1977), "Monopolistic competition and optimum product diversity", *The American Economic Review*, 67(3), pp.297-308.

Draganska, M. and D. Jain (2005), "Product-line lenght decisions as a competitive tool", *Journal of Economics and Management Studies*, 14, pp.1-28.

Dunne T, Klimex S, Roberts M and Y Xu (2005), "Entry and exit in geographic markets", *mimeo*, Pennsylvania State University.

Eberly JC (1994), "Adjustment of consumers durables stocks - evidence from automobile purchases", *Journal of Political Economy*, 102(3), pp.403-36.

Ericson R and A Pakes (1995), "Markov-perfect industry dynamics - a framework for empirical work", *Review of Economic Studies*, 62(1), pp.53-82.

de Figueiredo John M and M Kyle (2006), "Surviving the gales of creative destruction: the determinants of product turnover", *Strategic Management Journal*, 27(3), pp.241-64.

de Fraja G. (1996), "Product line competition in vertically differentiated markets", *International Journal of Industrial Organization*, 14, pp. 389-414.

Foncel, J. and M. Ivaldi (2005), "Operating system prices in the home PC market", *The Journal of Industrial Economics*, 53(2), pp.265-97.

Fudenberg, D., Tirole, J. (1998), "Upgrades, tradeins, and buybacks", *RAND Journal of Economics*, 29, pp. 235–258.

Genakos, C. (2004) "Differential merger effects: the case of the personal computer industry", Economics of Industry Papers, EI/37, Suntory and Toyota International Centres for Economics and Related Disciplines, LSE, STICERD.

Genakos (2005), "Product life cycles in the personal computer industry", *mimeo*, Cambridge University.

Geske, M.J., Ramey, V.A. and M.D. Shapiro (2004) "Why do computers depreciate?", NBER Working Papers 10831.

Gordon, B. (2006), "A dynamic model of consumer replacement cycles in the PC processor industry", *mimeo*, Carnegie Mellon University.

Gotz G. and K. Gugler (2004), "Market concentration and product variety under spatial competition: evidence from retail gasoline", CESifo Working Paper No. 1289, CESifo GmbH.

Gowrisankaran G. and M. Rysman (2006), "Dynamics of consumer demand for new durable goods", *mimeo.*

Gowrisankaran G and RJ Town (1997), "Dynamic equilibrium in the hospital industry", *Journal of Economics and Management Strategy*, 6(1), pp.45-74.

Greenstein SM and JB Wade (1998), "The product life cycle in the commercial mainframe market, 1968-1982", *Rand Journal of Economics*, 29(4), pp.772-89.

Kumar P. (2002), "Price and quality discrimination in durable goods monopoly with resale trading", *International Journal of Industrial Organization*, 20, pp.1313-39.

Johnson J. and D. Myatt (2003), "Multiproduct quality competition: fighting brands and product line pruning", *American Economic Review*, 93(3), pp.748-74.

Johnson J. and D. Myatt (2006), "Multiproduct oligopolists", *Rand Journal of Economics*, 37(3), pp.583-601.

Judd K. (1985), "Credible spatial preemption", *Rand Journal of Economics*, 16, pp.153-166.

Hausman J. Leonard G. and J. Zona (1994), "Competitive analysis with differentiated products", *Annales d'Economique et de Statistique*, 34, pp.159-80.

Hausman J. and Wise (1978), "A conditional probit model for qualitative choice: discrete decisions recognizing interdependence and heterogeneous preferences", *Econometrica*, 46(2), pp.403-26.

Hendel I. and A. Lizzeri (1999), "Interfering with secondary markets", *RAND Journal of Economics*, 30(1), pp.1–21.

Hendel I. and A. Nevo (2006), "Measuring the implications of sales and consumer stockpiling", *Econometrica*, 74(6), pp.1637-1673.

Hotelling, H. (1929), "Stability in Competition," *Economic Journal*, 39, pp.41-57.

Lam P. S. (1991), "Permanent income, liquidity and adjustment of automobile stocks: evidence from panel data", *Quarterly Journal of Economics*, 106, pp.203-30.

Lee I. and J. Lee (1998), "A theory of economic obsolescence", *The Journal* *of Industrial Economics*, XLVI(3), pp.383-401.

Levinthal D.A. and D. Purhoit (1989), "Durable goods and product obsolescence", *Marketing Science*, 8, pp.35-56.

Mazzeo, M. J. (2002), "Product choice and oligopoly market structure", *RAND Journal of Economics*, 33, pp.221-42.

McFadden, D. (1981), "Econometric model of probabilistic choice", in C.F. Manski and D.McFadden eds., *Structural analysis of discrete data*, MIT Press.

Melnikov, O. (2000): "Demand for differentiated durable products: The case of the U.S. printer market", mimeo, Yale University.

Nevo A. (2000), "Mergers with differentiated products: the case of ready-to-eat cereal indsutry", *Rand Journal of Economics*, 31(3), pp. 395-421.

Nevo A. (2001), "Measuring market power in the ready-to-eat cereal industry", *Econometrica*, 69(2), pp.307-42.

Pakes, A. (2003). "A reconsideration of hedonic price indices with an application to PC's", *American Economic Review*, 93(5), pp.1578-96.

Pakes A and P McGuire (1994), "Computing markov-perfectr nash equilibria: numerical implications of a dynamic differentiated product model", *Rand Journal of Economics*, 25(4), pp.555-89.

Pakes A, Ostrovsky M and S Berry (2004), "Simple estimators for the parameters of discrete dynamic games (with entry/exit examples)", *NBER Working Paper*, No. 10506.

Porter R. and P. Sattler (1999), "Patterns of trade in the market for used durables: theory and evidence", *NBER Working Papers*, No. 7149.

Prince, J. (2007) "The Beginning of Online/Retail Competition and Its Origins: An Application to Personal Computers", *International Journal of Industrial Organization*, 25(1), pp.139-56.

Requena F. and J. Walker (2003), "Competition and survival in differentiated product markets: the case of the UK automobile market (1971-1998)", mimeo.

Ryan S (2006), "The cost of regulation in a concentrated industry", *mimeo*, MIT.

Rust, J. (1986), "When is it optimal to kill off the market for used durable goods?" *Econometrica*, 54, pp.65–86.

Seim, K. (2006), "An empirical model of firm entry with endogenous product-type choices ", *RAND Journal of Economics*, 37(3), pp.319-40.

Siebert R. (2003), "The introduction of new product qualities by incumbent firms: market proliferation versus cannibalization", WZB Markets and Political Economy Working Paper No. SP II 2003-11.

Small, K.A. and S. Rosen (1981), "Applied welfare economics with discrete choice models", *Econometrica*, 49, pp.103-30.

Song M. (2007), "Measuring consumer welfare in the CPU market: an application of the pure characteristics demand model", *Rand Journal of Economics*, forthcoming.

Stavins J. (1995), "Model entry and exit in a differentiated-product industry: the personal computer market", *The Review of Economics and Statistics*, 77(4), pp.571-84.

Stavins, J. (1997) "Estimating Demand Elasticities in a Differentiated Product Industry: The Personal Computer Market", *Journal of Economics and Business*, 49(4), pp. 347-67.

Stengos, T. E. Zacharias (2006) "Intertemporal Pricing and Price Discrimination: A Semiparametric Hedonic Analysis of the Personal Computer Market" *Journal of Applied Econometrics*, 21(3), pp. 371-86.

Shaked, A., and Sutton, J. (1982), "Relaxing price competition through product differentiation", *Review of Economic Studies*, 49 3-13.

Shaked, A., and Sutton, J (1983), "Natural Oligopolies", *Econometrica* 51, pp. 1469-84.

Shaked, A., and Sutton, J.(1987), "Product Differentiation and Industrial Structure", *Journal of Industrial Economics*, 36, pp. 131-146.

Spence, M. (1976a), "Product Selection, Fixed Costs and Monopolistic Competition", *Review of Economic Studies,* 43, pp. 217-35.

Spence, M. (1976b), "Product Differentiation and Welfare" *American Economic Review*, 66, pp. 407-14.

Sutton J. (2001), "Rich trades, scarce capabilities", STICERD discussion paper EI/28.

Swan, P. (1970), "Durability of Consumption Goods." *American Economic Review*, 60, pp. 884–894.

Swan P. (1972), "Optimum durability, second hand markets, and planned obsolescence", *Journal of Political Economy*,80, pp.575-85.

Swan P. (1980), "Influence of recycling on monopoly power", *Journal of Political Economy*, 88(1), pp.76-99

Waldman, M. (1996), "Eliminating the market for secondhand goods: an alternative explanation for leasing," *Journal of Law and Economics*, 40, pp.61-92.

Watson, R. (2004), "Product variety and competition in the retail market for eyeglasses", Working Paper 0047, The Center for the Study of Industrial Organization, Northwestern University.