

UC Berkeley

UC Berkeley Electronic Theses and Dissertations

Title

On the Roles of Information in the Interaction of Agents in Markets

Permalink

<https://escholarship.org/uc/item/7jq5q5cw>

Author

Gardete, Pedro

Publication Date

2011

Peer reviewed|Thesis/dissertation

On the Roles of Information in the Interaction of Agents in Markets

By

Pedro Miguel Gardete

A dissertation submitted in partial satisfaction of the
requirements for the degree of
Doctor of Philosophy

in

Business Administration

in the

Graduate Division
of the
University of California, Berkeley

Committee in Charge:

Professor J. Miguel Villas-Boas, Chair

Professor Ganesh Iyer

Professor John Morgan

Professor Denis Nekipelov

Spring 2011

Copyright 2011

by

Pedro Miguel Gardete

Abstract

On the Roles of Information in the Interaction of Agents in Markets

by

Pedro Miguel Gardete

Doctor of Philosophy in Business Administration

University of California, Berkeley

Professor J. Miguel Villas-Boas, Chair

Markets are rich environments for study. Two features in particular make them attractive settings for the study of the roles of information. First, they are often characterized by a significant amount of uncertainty, an essential ingredient for information to be relevant. Second, strategic interactions abound.

Each chapter of this dissertation uses a different ‘lens’ so as to focus on particular roles of information in strategic market settings. First, information is considered in its pure ‘aid to decision-making’ role, in which agents use it to improve their decisions, not only based on the expected future conditions but also on the expected actions of the other agents which in turn depend on the information they are likely to hold. Second, information is considered in its ‘influencer’ and ‘coordinator’ roles. It can be used to influence the decisions of agents which will in turn affect the welfare of the original transmitters, while increasing overall market coordination. Finally, information is considered in its ‘homogenizer’ role, in the sense that actions and preferences may become more similar after its repeated use and subsequent transmission. The purpose of this study is to propose, document and estimate mechanisms by which information fulfills these specific roles in actual market settings.

This dissertation is organized in the following way. Chapter 1 analyzes the use of information acquired by decision-makers in a strategic context. Chapter 2 introduces the layer of the information transmitter as a separate strategic entity, who may have different interests from those of the receivers. Finally, Chapter 3 considers the case where the same agents are simultaneously receivers, users and transmitters of information.

These scenarios are analyzed in well-defined market contexts. Chapter 1 looks at the role of information in the production and capacity decisions of DRAM producers. I estimate the precision of market information used by firms in the DRAM industry, implicit in the patterns and accuracy of their actions. The value of information is estimated as well as its impact on the decisions of the firms.

Chapter 2 investigates advertising between firms and consumers in a market for a search good. It is shown that even if firms have an interest in providing biased information through

advertising, it is still possible for it to convey credible content that may be useful to consumers. In addition it is shown that the precision of information delivered is sparser as the quality of the product or service in question decreases.

Finally, Chapter 3 analyzes the context of online movie ratings and estimates how the opinions of consumers shape the opinions of subsequent consumers, *ad infinitum*. It is estimated that roughly half of the consumers' valuations for movies is a function of the valuations of other consumers who have had a chance to rate the movie earlier. A sensitivity analysis of how some extreme initial ratings influence consumer valuations is performed, revealing considerable effects.

To dad.

Contents

List of Figures	v
List of Tables	vi
Acknowledgements	vii
1 Information as an Aid to Decision-Making in a Strategic Setting: The case of the DRAM Industry	1
1.1 Introduction	1
1.2 Description of the DRAM Market	2
1.2.1 A Race for Memory Density	3
1.2.2 International Competition	3
1.2.3 Turbulent Prices Descents	4
1.3 Related Literature	5
1.4 Data	6
1.5 Revisiting the Static Competition Model	8
1.5.1 Demand and Information Structure	8
1.5.2 Firm Decisions	9
1.6 Dynamic Model for the DRAM Market	11
1.6.1 Production Decisions	11
1.6.2 Capacity decisions	12
1.6.3 State Variable Transitions	14
1.7 Estimation	14
1.7.1 Demand Parameters	14
1.7.2 Static Parameters	15
1.7.3 Dynamic Parameters	18
1.8 Counterfactual Analysis	19
1.8.1 Information Structure and Production Decisions	20
1.9 Conclusion	20
1.10 Figures	21

2	Information in its Influencer and Coordinator Roles: Advertising through Information	28
2.1	Introduction	28
2.2	Model	31
2.2.1	Equilibrium Concept and its Implications	34
2.2.2	Purchase and Search Decisions	35
2.2.3	Pricing Decision and Comparative Statics	36
2.2.4	Demand Heterogeneity and Consumer Surplus	38
2.3	Advertising Strategies and Market Outcome	40
2.3.1	Matching Interests	42
2.3.2	Incentive to Misreport	43
2.4	Welfare Implications	44
2.5	Extensions	45
2.5.1	Costly Advertising	45
2.5.1.1	Matching Interests	45
2.5.1.2	Incentive to Misreport	46
2.6	Conclusion	46
3	Information as a Homogenizing Force: The Impact of Others' Opinions in Online Ratings	48
3.1	Introduction	48
3.2	Data	52
3.3	The Model	54
3.3.1	The Gibbs Sampler	56
3.3.2	Simulation	56
3.4	Estimation Results	57
3.5	Sensitivity analysis	60
3.6	Conclusion	62
	Bibliography	64
A	Appendix to Chapter 1	70
A.1	Posterior Distribution Given Demand Signal	70
A.2	Polynomial approximation of the Value Function	70
A.3	First-Order Condition for Capacity Decisions	71
A.4	Dynamic Parameter Estimates	71
B	Appendix to Chapter 2	72
B.1	Quality Range	72
B.2	Interior Profit Maximization	73
B.3	Proof of Proposition 1	75

B.4	Proof of Proposition 2	75
B.5	Proof of Proposition 3	76
B.6	Definition of $q_{(1)}$ to $q_{(3)}$	76
B.7	Equilibria with Matching Interests	76
B.8	Incentive to Misreport	77
B.9	Continuous Quality	78
C	Appendix to Chapter 3	80
C.1	Generating Y	80
C.2	Generating $\varepsilon_{,0}$	81
C.3	Generating α	82
C.4	Generating Φ	83
C.5	Generating σ^2	84
C.6	Generating K	84

List of Figures

1.1	Timing	11
1.2	A 12-inch (300 mm) Silicon Wafer	21
1.3	Optimal Simulated Quantities given Demand Signals with and without Capacity Constraints	22
1.4	Simulated Effects of the Variance of the Demand Signal	23
1.5	Estimated Demand Curve and Respective 25%/75% Error Quantiles	24
1.6	Estimated Demand Shocks and Demand Signals for each Alliance (Dollars per GB)	25
1.7	Operating Profits by Alliance in three distinct information scenarios (Millions of Dollars)	26
1.8	Total Market Output Levels and Model Prediction	27
2.1	Timing	35
2.2	Ex-ante consumer surplus	40
2.3	ex-ante expected consumer surplus when facing price p	41
3.1	Average ratings changing over time for three randomly chosen movies. . . .	53
3.2	Distribution of ratings for 5 randomly chosen movies.	53
3.3	Convergence plots using simulated data for (a) α_1 (b) σ^2 (c) k_2 and (d) ϕ . .	57
3.4	Convergence plots using simulated data for (a) α_1 (b) σ^2 (c) k_2 and (d) ϕ for movie id 662.	58
3.5	Simulated ratings showing the effect of incorporating cumulative average rating in evaluations.	61
3.6	Simulated ratings showing the effect of early shock on subsequent evaluations.	62

List of Tables

1.1	Shift in Geographical Distribution of DRAM Production	4
1.3	DRAM Firms in Sample, Nationalities and Alliances	7
1.4	DRAM Demand Parameter Estimates	15
1.6	DRAM Information and Cost Parameter Estimates	17
1.8	Dynamic Parameter Estimates	19
2.1	Ex-ante consumer surplus for different advertising policies	45
3.1	Posterior means and standard errors for Model 2 using three different Gibbs chains.	59
3.2	Posterior mean and standard error of fixed effects for two models.	60
A.1	Dynamic Parameter Estimates	71

Acknowledgments

I would like to start out by acknowledging professor J. Miguel Villas-Boas, my advisor at the Haas School of Business. Throughout the last six years I have accumulated a steep debt of personal and professional advice as well as encouragement. I can only promise to ‘pay it forward’ to any student who falls under my guidance, as best as I can. Thanks. Thanks. Thanks. I simply cannot imagine a better advising style, so well put into practice.

I would also like to thank professors Ganesh Iyer, John Morgan and Denis Nekipelov for extending their time and *fides* over my projects and for having influenced them so decisively. Thank you.

I must also thank other professors who took their own time to interact about research ideas: Eduardo B. Andrade, Hai Che, Ying-Ju Chen, Clayton Critcher, Richard Gilbert, Benjamin Hermalin, Teck Ho, Ming Hsu, Shachar Kariv, Zsolt Katona, Botond Koszegi, Leif Nelson, Minjung Park and Matthew Rabin. All of them advised me on parts of the dissertation and made my work better.

Before this dissertation started to be imagined and written I was recommended by professors Fátima Barros, Fernando Branco, Pedro Oliveira and Francisco Veloso for the PhD program. I extend my thanks for their trust and support.

My colleagues and friends have left an indelible mark on my work and took the time to help me with various research matters: Yakov Bart, Ron Berman, Mario Capizzani, Bianca Cerchiai, Debbie Chiou, Albert Buixade Farré, Deepak Hegde, TI Kim, Chan Jean Lee, James Sawhill, Diogo Silva, Shubhranshu Singh and Kaifu Zhang.

A PhD entails a lot of courage and determination. The friends I met in Berkeley as well as my good friends from Portugal steered me in the right way during this endeavor. Thank you.

The decision to undergo a doctoral program was only possible with the blessing of my family. A special note of thanks goes to my mom, who had to ‘let go’ in quite an abrupt way for Portuguese standards. I also thank my whole family for cheering me on.

Finally, because tangible means are also required to finish a PhD, I would like to thank the indispensable support of the Pinto-Fialon graduate fellowship as well as the support of the Fundação para a Ciência e Tecnologia through the Programa Operacional Potencial Humano (POPH)/Fundo Social Europeu (FSE), partially funded by the European Union. I acknowledge their support. I include a graphic acknowledgment below.

I would also like to thank donors of the Haas School of Business, who helped me considerably: I intend to imitate your kind gestures in the future.



– Go Bears!

Chapter 1

Information as an Aid to Decision-Making in a Strategic Setting: The case of the DRAM Industry

1.1 Introduction

Consider the following sequence of events: on August 15th 2007, the EETimes reported that the DRAM (Dynamic Random Access Memory) industry had reasons “for cautious optimism”, mostly because of signs of stronger demand ahead, driven by sales of handsets and game consoles. Two weeks later, iSuppli revised its own market forecast with a \$2.5Bn cut in annual revenue, but announced a prediction of a 17.5% growth in sales for 2008. By mid-2008, Semiconductor Industry Association president George Scalise reported that DRAM revenues had declined by 34% in the first four months of 2008.¹

The DRAM industry is only one of many industries who depend heavily on assessments of the economic context to make a profit. Competitors face a number of decisions - including production orders and capacity investments - that are taken before a significant amount of uncertainty is resolved. The lag that spans from deciding capital investments to their full implementation is of approximately six months, while the throughput time of units in the production process is of approximately three months. These lag times in conjunction with high demand uncertainty create a turbulent setting for all the firms in the industry.² This paper analyzes how the industry dynamics depend on the information held by firms in a competition setting. It takes into account the fact that managers rely on public and private

¹ Sources: www.marketwatch.com/story/pressures-mounting-anew-on-dram-chip-market-isuppli-says, www.eetimes.com/showArticle.jhtml?articleID=201800294, and www.goo.gl/QWTh.

² The October 19th 2009 edition of the Wall Street Journal features a piece describing the volatility of the DRAM industry with an illustrative title: “The Chip Sector’s Boom or Bust DRAMa”.

information to make decisions, and that they expect their competitors to do the same based on the information they are likely to hold.

Because accessing reliable and complete data on public and private information is infeasible, I instead follow the motto *facta non verba* (or its english equivalent, actions speak louder than words) and recover the *unobservable* public and private information that is consistent with the *observable* data on firms' actions by use of a game-theoretic model. The model used for analysis contemplates decisions that carry implications for the present and the future. Production decisions affect the profit levels according to the true realization of demand. Capacity decisions carry significant dynamic effects since they determine the competitive conditions for the following periods, not only by limiting future production quantities but in addition by influencing the future investment costs of the firms.

After estimating an econometric model from first-principles it is possible to construct different scenarios and describe how the industry is affected by them. These scenarios allow us to calculate the value and precision of the information in the industry as well as to consider the case of information-sharing.

The specificities of the DRAM market are discussed next. Section 1.3 reviews the related literature. The dataset is described in Section 1.4. Section 1.5 revisits the static duopoly Cournot model under demand uncertainty. Section 1.6 describes the dynamic model and Section 1.7 presents the estimation procedure and the empirical results. Section 1.8 describes counterfactual analyses and Section 1.9 concludes.

1.2 Description of the DRAM Market

DRAM is a technological variant of what is generically called RAM or 'computer memory'. It is used in a number of electronic appliances like computers, cell phones, printers, game consoles among others. Intel is generally credited for having introduced the first DRAM chip in 1970: the 1103 model featuring 1Kb of storage capacity - 1024 bits - which competed with the dominant technology until then, the magnetic core. Besides Intel, two other American companies started DRAM production in the same year with similar products: Mostek and Advanced Memory Systems (the latter effectively entering the market before Intel, but with a less popular product).³ The success of DRAM over its predecessor technology was primarily due to its compactness which translated into less space required for each bit of storage, and in addition due to its faster read/write performance. DRAM dominated the volatile memory industry in less than a decade. A number of fundamental characteristics of the market were soon revealed and still survive until today. They are discussed below in turn.

³See Murillo (1993) for a description of the DRAM industry from 1970 until 1993.

1.2.1 A Race for Memory Density

The production of DRAM chips requires hundreds of procedural steps, and has been described as “one of the most difficult processes of high technology known to man” (Murillo, 1993). This process has only gotten more complex over time, as there has been a continued effort to introduce more memory into the same amount of physical space. Memory chips are etched on a silicon disc called wafer (see Figure 1.2 for a picture of one). Given that the cost of a wafer is relatively constant across production technologies, it follows that fitting double as many memory units into the same wafer results in roughly halving the variable production costs. Memory density is the fundamental cost determinant for the DRAM production process. Production decisions depend strongly on it, over and above other cost factors such as labor wages or the price of silicon and other raw materials.

Production yield - the usable number of chips per silicon wafer - is a key success factor in the industry. The 1103 chip's early market leadership occurred partly because of Intel's achievement of a 10% die yield, while at the same time competitors were known to achieve yields of 3% to 6%. As of 1997, a UC Berkeley study showed that the typical average line yield (the percentage of wafers that were not scrapped from the production process) was of 93%, and the die yield was of 77.4%, resulting in a ~72% effective chip yield for 200mm wafers (see ICE Corporation, 1997, Chp. 3).

1.2.2 International Competition

Except for its earlier stages, competition in the DRAM industry has taken place at the international scale. The introduction of the 1103 chip in 1970 had already caused the appearance of clones by other 19 companies as of 1972, including the first Japanese rival, NEC. By 1976 two Japanese companies had subsidiaries in the U.S.A., and one of them (NEC) was receiving its first big order of DRAM chips from Honeywell (\$5 million worth). This was the fruit of continued efforts by the Japanese companies to address reliability and provide fast access performance in their chips. Soon, the so-called ‘four sisters’ (NEC, Toshiba, Fujitsu and Hitachi) had managed to effectively enter the American market.⁴

In October 1985 Intel announced it was leaving the DRAM production altogether in order to focus on the microprocessor, a bet that proved successful *a posteriori*. In the same year Motorola declared it too was abandoning the market.⁵ A period of Japanese dominance in the DRAM market had started, even if short-lived. As of September 1986, a U.S.A.-Japan Semiconductor Agreement forced Japanese firms to stop “dumping” across world markets, and established “fair market values” for their products. The agreement was aimed at benefitting American and European producers through raising Japanese product

⁴A number of enabling factors are documented: availability of qualified human resources, low cost of capital, “deep pockets” linked to the Keiretsu firm structure, among others. By 1980 the differences in product reliability were measured repeatedly, giving significant advantage to the Japanese producers.

⁵Both firms reentered the market in 1987 as resellers.

Table 1.1: Shift in Geographical Distribution of DRAM Production

Region	Number of DRAM Makers		Market Shares	
	1992	2009	1992*	2009
US	3	1	5-8%	11.0%
Japan	12	1	65-70%	2.4%
Korea	3	2	22-25%	50.6%
Europe	2	1	2%	9.6%
Taiwan	—	3	—	8.1%

Distribution of DRAM producers by geography, and world unit market shares by the end of 1992 and in February 2009 (only main firms, amounting to 95% market share). * - 1992 shares refer to the 1Mb chip market. Source: Murillo (1993), and Gartner as cited at fabtech.com.

prices. Instead, it changed the competition landscape in an unpredicted fashion: it opened the door for Korean firms who had been acquiring technology licenses and had started to participate in the capital of American companies (e.g.: Samsung purchased 2.7% capital of the American DRAM manufacturer Micron in the mid-80's).

Concurrently, the quality-orientation that had served Japanese firms well in the past may have played a role in their demise. The shift in use of DRAM from mainframes to personal computers privileged the quest for low costs instead of quality and reliability.⁶ Yunogami (2006) argues that the quality-centric culture of Japanese firms hindered their ability to adapt to the new cost trend, and gave way to the Korean and Taiwanese newcomers. Korean firms eventually surpassed both Japan and the United States in production quantities, as described in Table 1.1. Currently, the DRAM market is organized into roughly more than 10 players, many of which are organized into alliances of firms which effectively share production technologies and coordinate R&D efforts, decide production levels and allocate available capacity by its members.

1.2.3 Turbulent Prices Descents

Price descents have been part of the DRAM market since its beginning. For instance, the price of the 1103 chip came down from \$60 in 1970 to only \$4 three years later. The downward trend in prices per bit has remained until today. The reference spot price for a

⁶As an example, today's DRAM modules used in computer servers often have error checking and error-correcting mechanisms (e.g.: parity checking bits), which are absent in most of personal computers due to cost reasons.

Kilobit of DRAM was of 0.0003 cents of a dollar as of April 2010,⁷ a reduction of 7 orders of magnitude in price per memory unit with respect to the original level of 1970. The improvements in production technologies are the main factor responsible for long-term decreases in market prices. Similar to what happens in other semiconductor industries, the price per unit decline is a result of firm coordination around Moore's law. In the DRAM industry, this law has been translated to state that the number of bits per wafer doubles every two years.⁸ Moore's law is today an effective focal point that firms use to coordinate the rhythm of increases in production efficiency, much in the spirit of Schelling (1960). It effectively drives prices down for semiconductor components over time.

Another empirical regularity is that the descent path is not smooth. Economic conditions impact the performance of the DRAM industry, pushing prices up and down around the descent path over time. This is a result of the lack of perfect information about demand by firms, in conjunction with a considerable lag between decision making and decision implementation.

1.3 Related Literature

There exist a number of theoretical contributions analyzing the role of information in market competition. Some examples are Novshek and Sonnenschein (1982), Clarke (1983), Vives (1984), Gal-Or (1985; 1986) and Shapiro (1986). They analyze the roles of information and the incentives of firms to share it in market settings.⁹ The most prevalent finding is that firms benefit from more accurate own private information, be it about levels of future demand or about production costs. Sharing information about demand is found to be optimal under Bertrand competition but not under Cournot, while sharing information about costs is found to be optimal under Cournot competition, but not under Bertrand. This relates to how information sharing changes the correlation of actions between the firms, and whether those actions are strategic complements or substitutes. For example, sharing good news about market demand will increase the correlation of production quantities. Because quantities are strategic substitutes under the linear demand specification, the net effect of information sharing is negative. On the other hand, when a firm shares private information about its low costs, it increases its production quantities and decreases the production of its rival. When firms face a linear demand, the negative correlation of quantities is beneficial for the firms as best-response functions are negatively sloped. The dual argument may be made for Bertrand competition.

⁷As noted in www.dramexchange.com on April 13th, 2010.

⁸See intel.com/pressroom/kits/events/moores_law_40th/.

⁹See also Raith (1996); Jin (2000) for a general approach to information sharing in markets. The role of market information in competition contexts has also been analyzed in the context of distribution channels (e.g.: Chu, 1992; Lariviere and Padmanabhan, 1997) and of specific settings of information sharing (e.g.: Villas-Boas, 1994; Chen et al., 2001).

The analysis of the role of information and its characteristics in strategic settings has been an almost exclusive domain of research in auctions.¹⁰ A noteworthy exception is provided by Armantier and Richard (2003) who analyze the impact of sharing cost information in the airline industry. They consider a static duopoly model where carriers can decide whether to operate in several markets and how much to ‘produce’ in each simultaneously. Sharing cost information significantly increases firms’ profits while it only moderately decreases consumer surplus overall. The current work extends that of Armantier and Richard (2003) by taking into account the role of market information in managerial decisions that have dynamic implications. While there exist several empirical contributions in dynamic market competition settings that include private information (see for example Pesendorfer and Schmidt-Dengler, 2003; Bajari et al., 2007; Aguirregabiria and Mira, 2007; Bajari et al., 2009; Ryan, 2009; Fershtman and Pakes, 2009), most of their focuses is in advancing estimation methods and none of them focuses on the role of information in markets. In addition, while these models include uncertainty in the cost structure, the DRAM industry (and semiconductor industries in general) provides a contrasting setting where the cost structure is predictable and fairly public, while demand is highly uncertain and plays a significant in the industry.

1.4 Data

The dataset has quarterly information from the first quarter of 2006 until the third quarter of 2009. The period of analysis takes place after the complaints by US personal computer makers of collusive behavior by DRAM makers. These claims were proven to be true: they led to heavy penalties to DRAM companies in the years of 2003 to 2005. The executives of some of the most important firms in the industry - Samsung, Hynix, Elpida and Infineon - pleaded guilty to either price fixing or obstruction of justice in relation to a United States Justice Department’s probe on the matter. The panel used for estimation of the model was provided by a firm in the DRAM industry. It comprises detailed price, production and capacity information on six of the main firms in the DRAM industry who account for roughly 85% of the market share of the global DRAM market. These firms are organized into four alliances, which effectively behave as coordinated organisms (and hence are taken as individual decision-making units below).

In the beginning of the sample period and until the third quarter of 2008, ‘virtual’ capacity usage (not taking outages and maintenance stoppages into account) was at 100% and the DRAM market faced undersupply. The main problem for firms was to update their production capacity appropriately so as to be able to sell more units at lower costs. However, from the last quarter of 2008 until the end of the sample period capacity utilization

¹⁰See Athey and Haile (2006) for a literature review of empirical analyses of auctions.

Table 1.3: DRAM Firms in Sample, Nationalities and Alliances

Name	Nationality	Alliance Identifier
Samsung	South Korea	1
Hynix	South Korea	2
Powerchip	Taiwan	3
ProMOS	Taiwan	3
Elpida	Japan	3
Micron	USA	4

dropped dramatically. During that period marginal production costs finally played a direct role in the production levels of firms, which for some years had been set to equal the available capacity. Due to the market oversupply, firms refinanced heavily through either their parent companies or through renegotiations with banks in order to stay in operation. The medium-sized player Qimonda (not present in the sample data) was the single casualty after having filed for bankruptcy in the beginning of 2009. A listing of the firms included in the analysis as well as of their nationalities and partnerships is provided in Table 1.3.

While producing DRAM implies putting a different number of resources in place, it is by far the production technology that determines the relevant cost for decision-making about production. The cost of producing a ‘wafer’ is relatively constant across technologies. However, the memory density/efficiency of a wafer (i.e.: the number of bits per wafer) varies by production technology, and ultimately dictates the relevant production cost per unit (in our case measured in millions of gigabytes). These data are not published and are only accessible through expert analysis. For the purposes of this paper several meetings with marketing managers of a firm in the industry provided information on the density levels for the alliances in the sample. The experts described memory efficiency paths for each of the alliances. The efficiency data was provided in terms of indexes relative to the level of the most efficient firm in the beginning of the period, which was normalized to 1. These indexes are related over time through firm-specific transition parameters.

South Korean Samsung Electronics is the leader in market share throughout the sample period, capturing almost a third of the unit sales in the market. It is followed by Hynix which achieves more than 20% market share of DRAM unit sales. The remaining chunks of the market are shared among the medium size competitors (i.e.: Elpida and Micron with 19% and 13% market shares), and by smaller firms (Powerchip and ProMOS with less than 6% each).¹¹

While production increased for most firms along the sample period (and in all cases firms were producing more in mid-2009 than in the beginning of 2006), capacity usage

¹¹ See for example <http://www.isuppli.com/Memory-and-Storage/News/Pages/Taiwans-Powerchip-Surges-in-First-Quarter-DRAM-Ranking.aspx>.

dropped from full utilization (which occurs until the third quarter of 2008) to roughly 80% until the third quarter of 2009. The discrepancy between productive capacities and unit productions depends on both capacity and production decisions by firms. These in turn, depend on two sets of factors: capacity and production costs on the one hand, and on market information on the other. It is therefore important to incorporate both these sets of factors into the game theoretic model.

1.5 Revisiting the Static Competition Model

This section revisits the “static” duopoly model of competition analyzed by the theoretical work mentioned above, and is useful to introduce the more complex model of the following section. It considers a linear and a constant elasticity specification for demand. Exogenous production capacities are introduced. The timing of the one-shot competition model is as follows: firms receive a private demand signal s_i , $i = 1, 2$. They then decide how much to produce, taking into account that their signal may be informative about other firms’ actions. Finally, the true demand shock ε is realized and profits are allocated.

1.5.1 Demand and Information Structure

The linear and constant elasticity demand specifications are considered for a static simulation exercise. Both specifications have been widely used in the literature. They take the following forms:

$$P = \alpha + \beta Q + \varepsilon, \varepsilon \sim N(0, \sigma_\varepsilon^2) \quad (1.1)$$

and

$$P = \alpha Q^\beta + \varepsilon, \varepsilon \sim N(0, \sigma_\varepsilon^2) \quad (1.2)$$

where P is the market price firms will sell the total market production Q . ε is an i.i.d. normally distributed error term that reflects the total market uncertainty under absence of information.¹²

Firms hold information about the demand shock in the form of a signal $s_i = f(\varepsilon, \eta, \xi_i)$ centered at the true demand shock, i.e., $E[s] = E[\varepsilon]$. The demand signal summarizes private information as well as managerial assessments and interpretations of available data and of the market outlook. The η component captures industry-wide biases, while ξ_i captures firm

¹²The assumption of unbounded support for the error term implies a probability that the price may be negative. Using the alternative constant elastic specification $\log(P) = \alpha + \beta \log(Q) + \varepsilon$, $\varepsilon \sim N(0, \sigma_\varepsilon^2)$ would avoid this problem as it relates the normally-distributed error term to the variance of the logarithm of price. However, this introduces non-linearities that would prevent us from interpreting the variance of the error term and of the demand signals exclusively as price uncertainty, since they would impact the price levels themselves.

idiosyncratic information. Hence, on a given period all firms may have biased information with respect to the true demand shock.

The following additive information structure is assumed:

$$s_i = \varepsilon + \eta + \xi_i \quad (1.3)$$

with $\eta \sim N(0, \sigma_\eta^2)$ and $\xi_i \sim N(0, \sigma_\xi^2)$. This induces the joint conditional probability function:

$$\begin{matrix} s_i \\ s_j \end{matrix} \bigg| \varepsilon \sim N \left(\begin{bmatrix} \varepsilon \\ \varepsilon \end{bmatrix}, \begin{bmatrix} \sigma_\xi^2 + \sigma_\eta^2 & \sigma_\eta^2 \\ \sigma_\eta^2 & \sigma_\xi^2 + \sigma_\eta^2 \end{bmatrix} \right) \quad (1.4)$$

where the correlation between the firms' signals given the true demand realization is positive due to the common factor η . This structure captures potential correlations of demand signals, over and above those generated by the centeredness of the signals at the true demand shock ε .

1.5.2 Firm Decisions

Firms decide production levels before knowing the true demand realizations and consequently the market price that they will be able to attain. Because of strategic interaction, they take into account the decisions their competitors may make based on their own private information. Hence, firm i faces the objective function:

$$\begin{aligned} \max_{q_i} \quad & E_{\varepsilon, s_j} [P(q_i, q_j(s_j), \varepsilon) - c | s_i] q_i \\ \text{s.t.} \quad & q_i \leq K_i \end{aligned} \quad (1.5)$$

where $q_j(s_j)$ is the rival's production level dependent on her own private information. Parameter c denotes the common marginal cost, and K_i is firm i 's production capacity. This is a strategic-interaction problem with private information, much in the spirit of auction settings. Problem (1.5) is solved by the Kuhn-Tucker conditions, here depicted for firm i :

$$\frac{\partial E_{\varepsilon, s_j} [\pi_i(q_i, q_j(s_j), \varepsilon, c) | s_i]}{\partial q_i} - \lambda(s_i) = 0 \quad (1.6)$$

$$\begin{aligned} \lambda(s_i) (q_i - K_i) &= 0 \\ q_i &\leq K_i \end{aligned} \quad (1.7)$$

$$\lambda(s_i) \geq 0 \quad (1.8)$$

where $\pi_i(q_i, q_j)$ denotes the firm profit and function $\lambda(s_i)$ provides the shadow cost of capacity given demand signal s_i . These conditions allow for both corner and interior quantity solutions, depending on the parameter values. Finding the solution for this system entails solving the first-order conditions for any given signal for each firm. In order to do this, the system is solved at r potential values of the signals (which take place at different quantiles of the distribution of s_i , given the variance parameters). An approximation of the optimal strategies $q^*(s_i)$ for each firm is made through the use of flexible (Akima) splines, in order to accommodate the continuous support of the demand signals.¹³ The expected profit is approximated by a monomial quadrature of degree fifth evaluated at $2d^2 + 1$ points, where d is the number of integrals implied by the expectation operator (in effect it is equal to the number of firms). Hence, the number of equations implied by system (1.6) is of $2 \text{ firms} \times r \text{ points} \times 2 \text{ FOCs} = 4r \text{ equations}$. Constraints (1.7) and (1.8) are incorporated by penalizing overuse of production capacity, and by searching the correct space of the parameter values.


The optimal policies $q^*(s_i)$ for both demand specifications under presence and absence of capacity constraints are shown in Figure 1.3. They are linear on the demand signal for the linear demand specification, in line with the theoretical result by Radner (1962), and are nonlinear for the constant elasticity demand. Under these policies one can verify through simulation the effect of the information parameters σ_ξ^2 and σ_η^2 on the patterns of production and profits. In particular the role of σ_ξ^2 under the unconstrained case of the linear demand is shown in Figure 1.4 (the results for the constant elasticity demand and of changing σ_η^2 are identical). Figures 1.4 a) and b) plot the profit of firm i and the variance of its production on the precision of information parameter σ_ξ^2 respectively. Both are decreasing on σ_ξ^2 . This means that more accurate demand information is better for the firm (even if her rival is benefitting from better information herself). In addition, the variance of its quantity decreases as the quality of information decreases. This happens because the firm is more responsive to accurate rather than inaccurate information. Figures 1.4 c) and d) plot the covariance and the mean squared error of the quantities of the two firms in the market. As the quality of information decreases, the covariance of the firms' quantities decreases as well. On the other hand, the minimum squared error follows a non-linear pattern: when information is reliable, firms base their decisions on it; but because it is accurate for both firms, their productions tend to be close. On the other hand, as information is less accurate, symmetric firms are likely to receive different reports on demand and to produce different quantities depending on those data. Finally, when information is not accurate, firms ignore it and again start producing similar quantities in each period.

¹³For implementation I use $r = 7$. Variations in the number of signal levels had no significant impact on the simulations.

1.6 Dynamic Model for the DRAM Market

This section generalizes the static duopoly case to a more complete model which incorporates production and capacity decisions by firms, the two fundamental decisions in the DRAM industry. Because of resource fixity, production decisions are approximately taken one quarter in advance, while capacity decisions take two quarters to be implemented.¹⁴ These decisions are interdependent. Before they are made, firms take into account both short and long term demand information. Short term information refers to the demand level at period $t + 1$, and long-term information refers to demand in period $t + 2$, where each period refers to a quarter. Long-term signals for demand at time $t + 2$ are private at time t , and become public at time $t + 1$. Hence, at time t each firm has access to both public and private short-term demand information, and to private long-term demand information. The timing of the game is as follows:

Figure 1.1: Timing



1.6.1 Production Decisions

Firms face a static inverse demand curve for a homogeneous good

$$P_t = P(\theta, q_{it}, q_{-it}) + \varepsilon_t, \varepsilon_t \sim N(0, \sigma_\varepsilon^2) \quad (1.9)$$

where θ is a vector of demand parameters known to firms, and ε_t is a period-specific demand shock which firms know up to its distribution and the demand signals. The term q_{it} is the production quantity of firm i at time t , and q_{-it} is a vector of the production decisions of the remaining firms at time t . At period t , firm i (which will be used as a placeholder for

¹⁴These numbers are a result of conversations with industry professionals.

each firm unless otherwise noted) decides its production level for period $t + 1$ by solving the maximization problem

$$\begin{aligned} \max_{q_{it+1}} \quad & E_{\varepsilon_{t+1}, s_{-it+1}} [P(\theta, q_{it+1}, q_{-it+1}(s_{-it+1}, \bar{s}_{t+1})) + \varepsilon_{t+1} - c_{it} | s_{it+1}, \bar{s}_{t+1}] q_{it+1} \quad (1.10) \\ \text{s.t.} \quad & q_{it+1} \leq K_{it+1} \end{aligned}$$

where s_{it+1} is a private demand signal of firm i of demand at time $t + 1$ and $\bar{s}_{t+1}$ is a vector of public signals of demand at time $t + 1$. The short and long-term demand signals are uncorrelated over time and are distributed respectively according to:

$$\begin{aligned} \begin{matrix} s_{1t} \\ \vdots \\ s_{Nt} \end{matrix} \bigg| \varepsilon_t \sim N \left(\begin{bmatrix} \varepsilon_t \\ \vdots \\ \varepsilon_t \end{bmatrix}, \begin{bmatrix} \sigma_\xi^2 + \sigma_\eta^2 & \cdots & \sigma_\eta^2 \\ \vdots & \ddots & \vdots \\ \sigma_\eta^2 & \cdots & \sigma_\xi^2 + \sigma_\eta^2 \end{bmatrix} \right) \quad (1.11) \end{aligned}$$

$$\begin{aligned} \begin{matrix} \bar{s}_{1t} \\ \vdots \\ \bar{s}_{Nt} \end{matrix} \bigg| \varepsilon_t \sim N \left(\begin{bmatrix} \varepsilon_t \\ \vdots \\ \varepsilon_t \end{bmatrix}, \begin{bmatrix} \sigma_\xi^2 + \sigma_\eta^2 & \cdots & \sigma_\eta^2 \\ \vdots & \ddots & \vdots \\ \sigma_\eta^2 & \cdots & \sigma_\xi^2 + \sigma_\eta^2 \end{bmatrix} \right) \quad (1.12) \end{aligned}$$

such that the precision and the covariances of the short and long-term information is allowed to be different. Production costs and capacity levels at time $t + 1$ are common knowledge across firms. This is assumed to be the case for production costs because of the strong market coordination around Moore's law. Also, capacity levels for period $t + 1$ are known by all firms at time t because they are inferable and/or observable before being fully implemented. This is in line with the fact that information about capacity changes is often disseminated in the industry.

Finally, notice that production decisions only affect the next period's profits, but have not further consequences on the firm's payoffs. Capacity decisions on the other hand hold dynamic effects for the firms and have to be analyzed in terms of the sum of future payoffs.

1.6.2 Capacity decisions

Let $\pi_i^*(\Omega_t, \Theta_{it})$ denote firm i 's equilibrium operating profit at time t given the sets of state variables Ω_t and Θ_{it} . The first set $\Omega_t = \{c_t, \varepsilon_t, K_t\}$ contains a vector of marginal costs for all firms, c_t , the true demand shock for period t , ε_t , and the vector of capacity levels at time t , K_t . The set $\Theta_{it} = \{s_{it}, \bar{s}_t, \bar{s}_{it+1}\}$ contains the relevant market information about demand. These variables are relevant to characterize the dynamic problem. Since capacity levels constrain the future production levels and affect future investment decisions, firms need to take the discounted sum of future payoffs into account when making decisions over capacity levels. One way of writing those discounted payoffs is under the form of a Bellman equation, which relates the sum of the discounted cash-flows to today's financial impacts

and to the future payoffs under contingent decision-making. In our case the Bellman equation for firm i at time t can be expressed as

$$\begin{aligned}
W_i(\pi_i(\Omega_t, \Theta_{it}), \Omega_{t+1}, \Theta_{it+1}) = & \quad (1.13) \\
& \pi_i(\Omega_t, \Theta_{it}) + \max_{\substack{q_{it+1} \\ K_{it+2}}} \delta E_{\Omega_{t+2}, \Theta_{it+2}} [W_i(\pi_i(\Omega_{t+1}, \Theta_{it+1}), \Omega_{t+2}, \Theta_{it+2}) | \Omega_{t+1}, \Theta_{it+1}] \\
& - \text{InvCost}(K_{it+1}, K_{it+2}) \\
& \text{s.t. } q_{it+1} \leq K_{it+1}
\end{aligned}$$

where $W_i(\pi_i(\Omega_t, \Theta_{it}), \Omega_{t+1}, \Theta_{it+1})$ is the sum of discounted operational cash-flows at time t , for which the next period's costs, capacity levels and demand information are known. The value function depends on the expected sum of future operational profits and on today's capacity investment cost, $\text{InvCost}(K_{it+1}, K_{it+2})$. The investment cost is expected to be positive unless the scrap values for reducing capacity are high enough.

Finally, note that at time t firms choose both the production levels for time $t+1$ and the capacity levels for time $t+2$. The fact that production quantities only carry an effect into the next period enables us to separate production and capacity decisions by plugging in the equilibrium production quantities dependent on the state variables of program (1.10) into the Bellman equation (1.13), and rewriting the latter only as a function of the maximization problem over capacity levels. Let π_{it}^* denote the equilibrium operating profits of firm i at time t given the relevant state variables Ω_t, Θ_{it} . Bellman equation (1.13) becomes:

$$\begin{aligned}
W_i(\pi_{it}^*, \Omega_{t+1}, \Theta_{it+1}) = \pi_{it}^* + & \max_{K_{it+2}} \delta E_{\Omega_{t+2}, \Theta_{it+2}} [W_i(\pi_{it+1}^*, \Omega_{t+2}, \Theta_{it+2}) | \Omega_{t+1}, \Theta_{it+1}] \quad (1.14) \\
& - \text{InvCost}(K_{it+1}, K_{it+2})
\end{aligned}$$

Finally, note that firm i 's value function is separable into the next period's profits and the remaining payoffs. To see this define V_i such that

$$W_i(\pi_{it}^*, \Omega_{t+1}, \Theta_{it+1}) = \pi_{it}^* + V_i(\Omega_{t+1}, \Theta_{it+1}). \quad (1.15)$$

Together with expression (1.14), one gets the identity:

$$\begin{aligned}
V_i(\Omega_{t+1}, \Theta_{it+1}) = \delta \pi_{it+1}^* + & \max_{K_{it+2}} \delta E_{\Omega_{t+2}, \Theta_{it+2}} [V_i(\Omega_{t+2}, \Theta_{it+2}) | \Omega_{t+1}, \Theta_{it+1}] \quad (1.16) \\
& - \text{InvCost}(K_{it+1}, K_{it+2})
\end{aligned}$$

which relates the market outcome from period $t+1$ with the investment payoffs today and the remaining future cash-flows. Assuming there is an interior solution for capacity levels,

the equilibrium decisions must simultaneously satisfy the first-order conditions:

$$\frac{\partial}{\partial K_{it+2}} [\delta E_{\Omega_{t+2}, \Theta_{it+2}} [V_i(\Omega_{t+2}, \Theta_{it+2}) | \Omega_{t+1}, \Theta_{it+1}] - InvCost(K_{it+1}, K_{it+2})] = 0, \forall i \in N \quad (1.17)$$

1.6.3 State Variable Transitions

Because the demand shock is not serially correlated, demand signals have no bearing on demand shocks other the ones they relate to.¹⁵ On the contrary, production costs are decreasing over time by virtue of efforts of firms to accompany Moore's law, which in the DRAM context dictates that memory density per wafer should approximately double every two years. Hence, the memory density of firm i follows the process:

$$\tau_{it+1} = \phi_i \tau_{it} \quad (1.18)$$

where ε_{it} is firm i 's efficiency level at time t and ϕ_i denotes the firm-specific memory density transition factor. This rule translates directly into the transition rule for the unit production cost:

$$c_{it+1} = \phi_i^{-1} c_{it}. \quad (1.19)$$

where c_{it} is the marginal production cost of firm i at time t .

1.7 Estimation

1.7.1 Demand Parameters

DRAM is a significantly commoditized product and is sought by electronic product manufacturers. The following flexible demand specification is assumed:

$$P_t = \left(1 + \mu\lambda + \beta \left(Q_t^\lambda - 1\right)\right)^{1/\lambda} + \varepsilon_t, \varepsilon_t \sim N(0, \sigma_\varepsilon^2) \quad (1.20)$$

where P_t is the series of market prices and Q_t is the aggregate production at time t . This demand specification is based on the Box-Cox transformation of both dependent and independent variables (see Box and Cox, 1964) and includes the linear and the constant elastic demand curves as special cases. It reduces to the linear demand curve when λ equals one, and has the constant-elasticity specification as a limit of $\lambda \rightarrow 0$. The demand shock is

¹⁵A "log-log" constant elasticity specification was used to test error autocorrelation. A Prais-Winsten estimation revealed an autocorrelation parameter of 0.12 and a Durbin-Watson statistic of 1.78. Both statistics suggest absence of autocorrelation: this is an unusual finding in time-series estimation but in line with results of conversations with practitioners in the field.

Table 1.4: DRAM Demand Parameter Estimates

Parameter	Coefficient	Standard Error
μ	20.58	8.01
β	-0.777	0.160
λ	0.384	0.167
σ_{ε}^2	81.00	29.57*

* - Standard error was calculated from the maximum-likelihood information matrix. The market elasticity changed from -1.52 in the beginning of the period of analysis to -0.37 in the end of the sample period.

assumed to be i.i.d. normal with variance σ_{ε}^2 . The variance of the error term σ_{ε}^2 depends on the factors that affect the demand - such as economic conditions, shocks to producers of substitute and/or complement products, etc. - and is a measure of demand uncertainty under *full absence of market information*. Because firms have access to more information about demand than the econometrician and may use such information to make their decisions, it is important to correct for endogeneity. Although the use of input prices is a popular choice, they are mostly irrelevant in the DRAM industry, as discussed in the introduction. The main production factor is defined by the density technology in place, on which managers strongly base their production and capacity decisions on. Since memory density levels depend strongly on firms following Moore's law but are not related to demand shocks, they are valid instruments for the production quantities. The matrix of instruments Z is formed by a constant and the levels of memory densities. These are used to identify the demand parameters under the assumption $E[Z'\varepsilon] = 0$. Demand was estimated through an efficient 2-step generalized method of moments and the estimation results are presented in Table 1.4.

Figure 1.5 plots the average demand curve and willingnesses-to-pay scenarios at the 25th and 75th percentiles of the error term. In the third quarter of 2009 the revenue amplitude between these two quantiles was of \$726Mn, an eighth of the overall market size for that quarter, and should be even more in the remaining 50% of cases. However, firms hold information that partially shields them from this uncertainty.

1.7.2 Static Parameters

Throughout the sample period, and until the third quarter of 2008, there is a ramping-up in capacity levels. This is a result of ever-decreasing production costs allied to a demand that is able to absorb production at a low enough price. During this period the 'virtual' capacity usage (not taking outages and maintenance into account) was at 100%, and the DRAM

market faced undersupply. However, from the last quarter of 2008 until the end of the sample period capacity utilization dropped dramatically since DRAM producers could no longer get a profitable price for their production. In the latter periods marginal production costs finally played a direct role in production decisions of firms. For the purposes of this paper industry professionals were interviewed about the density levels for the firms in the sample across the data period. The cost per bit, c_{it} , can be expressed as:

$$c_{it} = \frac{\text{Cost per Processed Wafer}}{\text{Bits per Wafer}_{it}} \quad (1.21)$$

Let ε_{it} denote firm i 's efficiency index (memory density) at time t . In the beginning of the sample it is equal to 1 for the most efficient firm. Given that the same firm could fit 38.75 GB per each wafer at that time, equation (1.21) can be recast as:

$$c_{it} = \frac{\omega_0}{38.75\tau_{it}}$$

where ω_0 can be interpreted as the production cost related to acquiring and processing a memory wafer. The firm-specific transition factors for costs, ϕ_i , were provided by industry professionals in interviews. The estimation of the information and cost parameters relies on the static first order conditions of the profit optimization problem (1.10). They generate a similar system of equations as described for the duopoly case in (1.6), and are solved by the use of flexible spline functions approximating the firms' optimal strategies $q_{it}^*(s_{it+1}, \cdot)$ in each period, where s_{it+1} is the only relevant element of private information of each firm. For inference purposes the optimal quantity $q_{it}^*(s_{it+1}, \cdot)$ is allowed to go above the capacity level during the parameter search, in order to produce the correct moments of the distribution of the sample quantities. However, we would expect the predicted quantities provided by $q_{it}^*(s_{it+1}, \cdot)$ to approximate the actual quantities arbitrarily well as the sample size increases.

The conditional expectation of profits in (1.10) is approximated by a monomial quadrature of degree fifth over the posterior distribution of the demand shock and of the rivals' signals, $G_i(\varepsilon_t, s_{-it} | s_{it})$, evaluated at $2d^2 + 1$ points, where d is effectively equal to the number of firms. Static profit maximization hence implies a system of $4 \text{ firms} \times r \text{ points} \times 2 \text{ FOCs} = 8r \text{ equations}$, which is solved for each period in order to take changes of production costs and public demand signals into account.¹⁶

Let the generic vector $\{x_t^r\}_{r=1}^R$ denote a sequence of random variables x for time t , where r is an identifier of the simulation instance. The static first order conditions give rise to the predicted equilibrium quantities $\{\hat{q}_{it}(s_{it}^r, \cdot) | \hat{\varepsilon}_t\}_{t=1}^T$ for simulation r , where $\hat{\varepsilon}_t$ is the estimated demand shock from the first stage estimation. The predicted quantities are used

¹⁶For implementation I use $r = 7$. Variations in the number of signal levels had no significant impact on parameter estimates. Each evaluation of the GMM criterion function of the static estimation solves a system of 56 equations for each of the time periods.

Table 1.6: DRAM Information and Cost Parameter Estimates

Parameter	Coefficient	Standard Error
σ_{ε}^2	40.15	3.82
σ_{η}^2	109.32	7.67
ω_0	261.52	4.63

to form the simulated moments:

$$(TR)^{-1} \sum_{t=1}^T \sum_{r=1}^R (q_{it} - \hat{q}_{it}(s_{it}^r) | \hat{\varepsilon}_t), \forall i \in N. \quad (1.22)$$

Including the demand shock in this moment estimation is beneficial since demand signals are correlated with the true demand shocks. In addition, the orthogonality between the production costs and the demand signals and shocks gives rise to the additional moment equations:

$$(TR)^{-1} \sum_{t=1}^T \sum_{r=1}^R c_t (q_{it} - \hat{q}_{it}(s_{it}^r) | \hat{\varepsilon}_t), \forall i \in N. \quad (1.23)$$

Table 1.6 presents estimates of the information and of the cost parameters.¹⁷

The variance of the demand signals is lower than the variance of the error term. These results highlight the importance of correcting for demand endogeneity in demand estimation, since market information about demand seems to be relatively precise such that the firms' actions should be correlated with the demand shocks.

The posterior variance of the demand shock given the demand signal is given by

$$V(\varepsilon_{it} | s_{it}) = \left(\frac{1}{\sigma_{\varepsilon}^2} + \frac{1}{\sigma_{\varepsilon}^2 + \sigma_{\eta}^2} \right)^{-1} \quad (1.24)$$

and is equal to 38.51, a lower uncertainty faced by the firms than that implied by the unconditional variance of the demand shock. Figure 1.6 plots the recovered demand shocks and the demand signals of each firm alliance during the period of the last quarter of 2008 until the end of the sample. It is clear that alliances 1 to 3 overestimated the expected market conditions, while alliance 4 underestimated them. Given the size of alliances 1 to 3, it is clear that the market 'bust' during this period arises out of reliance in overoptimistic demand information. The fact that the conditional variance of the demand signal given the

¹⁷The parameters above were estimated based on the short-run demand information signals. The joint estimation of the static and dynamic moments is currently under way.

information is low makes firms rely on the information more. Although this is a winning strategy on average, it is bound to provoke mishaps when otherwise good information goes bad. Finally, the cost parameter reveals a wafer production cost of \$261.52. This value lies between the cost of a silicon wafer and the same cost added to the total processing costs. The estimated cost makes sense as it only reflects the variable production costs.

1.7.3 Dynamic Parameters

The dynamic parameters describe the value function and the investment costs of changing capacity levels. The investment cost function is assumed to be quadratic on changes in capacity levels:

$$InvCost(K_{it}, K_{it+1}) = [\gamma^+ 1(K_{it+1} > K_{it}) + \gamma^- 1(K_{it+1} < K_{it})] (K_{it+1} - K_{it})^2$$

where parameters γ^+ and γ^- capture the investment costs of positive and negative net changes in capacity levels, and $1(\cdot)$ is an indicator function. It is possible that γ^- becomes negative in case the scrap value of resources related to capacity is high enough.

Firm i 's value function is approximated by a polynomial function, $P_i(\Omega_{t+1}, \Theta_{it+1}) \simeq V_i(\Omega_{t+1}, \Theta_{it+1})$, in the spirit of Bajari et al., 2009.¹⁸ The value function identity and the first-order condition for capacity provide the moment equations for estimation of the dynamic parameters. To see this, consider the following simple example of a generic dynamic problem:

$$V(s) = \max_x \pi(s, x) + \delta E_{s'} [V(s') | s, x]$$

where s is a singleton state variable and x is a control variable that may influence the transition of the state variable s to s' . Let x^* be the optimal decision of the firm, which is observed by the econometrician. The Bellman equation can then be rewritten as

$$V(s) = \pi(s, x^*) + \delta E_{s'} [V(s') | s, x^*]$$

which, by moving elements to the left-hand-side, gives rise to:

$$V(s) - \pi(s, x^*) - \delta E [V(s') | s, x^*] = 0$$

Finally, because the econometrician assumes a flexible form for the value function and can observe the state variables (and in addition can use simulation to include directly unobservable variables), the expectation operator can be extended to the whole of the left-hand side:

$$E [V(s) - \pi(s, x^*) - \delta V(s') | s, x^*] = 0$$

¹⁸The expressions of the polynomial $P(\cdot)$ and of the first order conditions are presented in the appendix.

Table 1.8: Dynamic Parameter Estimates

Parameter	Coefficient	Std. Error
σ_ε^2	110.36	48.13
σ_η^2	200.5	21.53
γ^+	20.58	9.70
γ^-	8.26	5.05

which is the familiar conditional mean restriction that implies orthogonality conditions useful for the estimation. In the same spirit, the value function identity (1.16) provides the moments

$$((T-1)R)^{-1} \sum_{t=2}^T \sum_{r=1}^R \left[P_i(\widehat{\Omega}_{t+1}^*, \Theta_{it+1}^r) - (\delta\pi_{it+1}^* + \delta P_i(\widehat{\Omega}_{t+2}^*, \Theta_{it+2}^r), \forall i \in N \right. \\ \left. - \text{InvCost}(K_{it+1}^*, K_{it+2}^*) \cdot (\widehat{\Omega}_{t+1}^*, \Theta_{it+1}^r) \right] \quad (1.25)$$

where $\widehat{\Omega}_t^* = \{c_t, \hat{\varepsilon}_t, K_t^*\}$ and $\Theta_{it}^r = \{s_{it}^r, \bar{s}_t^r, \bar{s}_{it+1}^r\}$. The first order conditions of capacity decisions also provide moments for estimation:

$$((T-1)R)^{-1} \sum_{t=2}^T \sum_{r=1}^R \left[\frac{\partial}{\partial K_{it+2}} \left[\delta P_i(\widehat{\Omega}_{t+2}^*, \Theta_{it+2}^r) \right. \right. \\ \left. \left. - \text{InvCost}(K_{it+1}^*, K_{it+2}^*) \cdot (\widehat{\Omega}_{t+1}^*, \Theta_{it+1}^r) \right] \right] \Big|_{K_{it+2}=K_{it+2}^*}, \forall i \in N \quad (1.26)$$

The orthogonality of the moments with the lagged capacities were used, and in addition the production costs and the estimated demand error terms could also be added. The estimated parameters of the value function are presented in Table 1.8.

The variance of the long-term demand signals is higher than that of the short-run demand signals. In addition, the correlation of the long-term signals is of 84.4%, also higher than the correlation of short-term demand signals which is equal to 7% (and not significantly different from zero). Reducing capacity seems to be more costly than building it by approximately \$90,000 per million of gigabytes. The remaining parameters of the polynomial approximation of the value function are reported in the appendix in Table A.1.

1.8 Counterfactual Analysis

Once the model fundamentals are estimated, one can consider alternative scenarios. In particular the impact of different information structures on short-term decisions is considered.

1.8.1 Information Structure and Production Decisions

This section analyzes the effects of the information structure on the firms' profits, taking capacity decisions as given. The goal is to analyze the impact of the characteristics of market information on production decisions and the resulting outcomes. Two counterfactuals are considered: the case where perfect information is available to firms - which is generated by setting the variance parameters σ_{ξ}^2 and σ_{η}^2 equal to zero - and the case where no information is available - which is generated by setting the variance parameters σ_{ξ}^2 and σ_{η}^2 equal to infinity.¹⁹ This simulation is performed for the last quarter of 2008 onwards, when the last demand 'bust' took place. The comparison of firms' operating profits for that period in both cases and the actual history is described by Figure 1.7. As expected, perfect demand information dominates the remaining cases. In addition, in the last quarter of 2008 it would have been better for firms to have ignored the demand information rather than having used it. The overoptimistic market outlook led firms to overproduce and gave origin to the latest 'bust' event in the industry. High production levels and a low market price resulted, as seen in Figure 1.8. In particular in the period of 4Q 2008 until 3Q 2009 firms approximately halved their operational profits due to bad demand information in the end of 2008 and in the beginning of 2009. Finally, Figure 1.7 also shows that the value of information is positively correlated with the size of the firm, as Samsung's profits are the most sensitive (in absolute value) to variations in the quality of information.

1.9 Conclusion


Demand information is a fundamental component for decision making in competition settings. This is especially true in industries where decisions are made in advance to actual demand levels being known. This paper takes the example of the DRAM industry in order to estimate the characteristics of the market information available to managers and to analyze the structure of the quality of information on firms' profits.

As expected, short-run market information appears to be more accurate than long-term market information, and is also less correlated among firms. The costs of production and of changes in capacity levels were estimated, and the impact of the 'unexpected' low demand for the last quarter of 2008 until the end of the sample period were analyzed. In those four quarters demand information actually had a net negative impact as firms relied on overoptimistic forecasts for the end of 2008 until the beginning of 2009. The overall industry short-run operational loss due to inaccurate demand information is estimated to have been of approximately \$10Bn in operational profits.

¹⁹The analysis of the dynamic scenarios is currently under way.


1.10 Figures

Figure 1.2: A 12-inch (300 mm) Silicon Wafer




Silicon wafers are 'blank slates' for several semiconductor industries, like RAM, Flash memory, CPU's, etc.


Figure 1.3: Optimal Simulated Quantities given Demand Signals with and without Capacity Constraints




a)



b)




c)




d)

Underlying demand curve parameters are the ones estimated from the data. In addition: $c = 50.0$; $\sigma_\xi = 0.1$; $\sigma_\eta = 0.1$. a) and c) refer to the linear demand case, without and with capacity constraints, while b) and d) refer to the constant-elasticity demand case.


Figure 1.4: Simulated Effects of the Variance of the Demand Signal




a)



b)




c)



d)


Underlying demand curve parameters are the ones estimated from the data.
In addition: $c = 50.0$; $\sigma_\xi = 0.1$; $\sigma_\eta = 0.1$. The plots above are based on the average of 10,000 signals centered at different demand errors.

Figure 1.5: Estimated Demand Curve and Respective 25%/75% Error Quantiles




The dashed lines describe shifts in the demand function according to error shocks measured in the 25th and 75th percentiles.

Figure 1.6: Estimated Demand Shocks and Demand Signals for each Alliance (Dollars per GB)



The demand signals for the bigger alliances (1-3) predict higher market prices than the ones realized until the second quarter of 2009.


Figure 1.7: Operating Profits by Alliance in three distinct information scenarios (Millions of Dollars)



Note: All profit levels start at zero. During the downturn of late 2008/early 2009, overoptimistic demand information led firms to overproduce and deteriorate profits. As seen in the figure above, during this particular period it firms would have been better off by ignoring demand information overall.


Figure 1.8: Total Market Output Levels and Model Prediction

a)



The dashed line represents the prediction for market output given the true demand shocks over 1,000 simulated demand signals for each period. The positive trend is explained by decreasing unit cost reductions over time. The dashed box is depicted in below and describes the 2008 market slump.

b)



The dashed line represents the prediction for market output given the true demand shocks over 1,000 simulated demand signals for each period. The third quarter of 2008 is characterized by market overproduction relative to the optimal level implied by the estimated demand shock. This situation partially reverses in the following periods, and normalizes in the second quarter of 2009.

Chapter 2

Information in its Influencer and Coordinator Roles: Advertising through Information

2.1 Introduction

The previous chapter considered the case of a strategic setting where information is unbiased and non-strategically communicated to the decision-makers. I now focus on a setting where this assumption is relaxed. Consider the following scenario: Nielsen's April 2009 Global Online Consumer Survey reports that respondents have limited trust in firm advertisements. On a 0 to 100% scale TV and Radio ads scored 62% and 55% in consumer trust respectively, in contrast to recommendations from people known (90%) and consumer opinions posted online (70%).¹ In which conditions the content of such advertising messages is credible to consumers is the focus of this paper. A model of communication between a firm and consumers is analyzed, in which the firm can make claims about the quality of its product, which is its own private information. The model is applied to a market for a 'search good', in which the "*consumers' information problem is to evaluate the utility of each option*" - Nelson (1970). In this context, consumers observe the advertising message by the firm and decide whether to incur a search cost in order to learn more about the product, and potentially acquire it. The model predicts that informative communication is possible in two different cases. In the first, the firm has an incentive to report the quality of its product correctly since the interests of the consumers and those of the firm match. In the second, the firm may have an incentive to exaggerate its quality level despite it facing

¹Personal recommendations and consumer opinions posted online are the most trusted forms of advertising globally, Nielsen Newswire, http://blog.nielsen.com/nielsenwire/wp-content/uploads/2009/07/pr_global-study_07709.pdf

strategic consumers. In this case, the model describes a semi-separating equilibrium where information is sparser as product quality decreases.

In both scenarios the intuition for communication credibility is that different advertising messages attract and repel different consumer segments. Credibility arises because the firm makes trade-offs across consumer segments. For example, targeting a specific consumer with an ad about the high quality of a product may discourage other consumers because of the high price they will have to pay for it.

An important assumption in the current model is that different types of consumers have access to different outside options.² This idea is compatible with the interpretation of the marginal utility of quality (the ' θ ' parameter in some vertical differentiation models) being proportional to the inverse of the consumers' marginal utility for income (see Tirole, 1988, pp. 96-97). In that case, higher consumer types (in θ) hold higher incomes and may have access to different alternatives than those accessible to lower types. This assumption implies that only the firm holding a product of high enough quality prefers to serve the high end consumers, since it will be competing with a high outside option. It follows that whenever there is significant demand heterogeneity (i.e.: marginal valuations for quality differ significantly in the market), low type consumers may not always want to search if they believe that the quality of the product is very high. This is because they will anticipate that the firm selling such a product will be focusing on the high type consumers by setting a high price. In this sense, high quality is not always 'good news' for the lower type consumers.³ This setting gives rise to a matching problem between the firm and the consumers. The firm makes a strategic choice over its advertising message which may be truthful or not. Consumers are not required to believe the claim of the firm. Instead, they update their beliefs about the quality and price of the product and decide whether to search or not.

The present analysis describes how advertising can be credible because of its content, i.e., its message. It is different from the classical approach considered by Milgrom and Roberts (1986) where price or the observable cost of advertising serve as credible signals for consumers (see Kirmani and Rao, 2000 for a review of other credible signals in the context of adverse selection models in vertically differentiated models). In the current paper, advertising is first considered costless in order to isolate the informational power of the message, and later considered to be associated to a 'small' cost .

²The assumption that different types of consumers may have access to different outside options has been widely used in the literature. See Lewis and Sappington (1989); Maggi and Rodriguez-Clare (1995); Villas-Boas and Schmidt-Mohr (1999); Jullien (2000); Laffont and Martimort (2002) and Stole (2007).

³A similar case where apparent 'good news' can be bad for consumers in a search good setting is provided by Villas-Boas (2009), where a big product line (and hence a high probability of product fit) necessarily implies a higher expected price which can deter some consumers from searching.

The cheap-talk literature (mainly Crawford and Sobel, 1982) provides the underlying framework for costless (or cheap) communication between agents. The present work endogenizes the ‘sender bias’ presented by Crawford and Sobel (1982) by considering that firms may want to misreport their quality in order to attract different consumers and hence charge different prices in order to increase profits. Similar to Crawford and Sobel’s model, communication can be informative as long as the bias for miscommunication is not too high.

The advertising literature exhibits a small number of references on the provision of potentially biased information. Bagwell and Ramey (1993) initiate this work by proposing that in certain conditions, even though firms are able to misreport the true quality of their products, they may prefer to advertise truthfully if they share similar interests with consumers. The intuition for the existence of matching interests relates to free-entry and fixed costs. When fixed costs of entry of high and low quality firms are similar, misrepresentation does not alter their market shares by too much due to market entry. A low type firm claiming to be of high type cannot attract many more consumers by misreporting; it only attracts a different type of consumers. In that case there is no gain from misrepresentation: on the contrary, there would only exist the cost of attracting relatively the same number of consumers who would find the product unattractive relative to their preferred one, which will then generate less profits for the firm. In conclusion, when firms cannot attract a larger number of consumers through misrepresentation, they have an incentive to report their true product quality and consumers have an incentive to believe them. The problem I consider is the same in spirit as in Bagwell and Ramey (1993) although the mechanism I propose and the cases I analyze are somewhat different. In the current model, advertising can be informative because firms may have different preferences over types of consumers, even in a setting of a vertically differentiated market. The reason is because consumers who value quality more are the most likely to have access to higher outside options, as described above. In contrast, while allowing for their general result of informativeness through cheap-talk, I uncover cases in which communication is possible even though the interests of the firm and of the consumers do not match exactly (i.e.: incentive to misreport exists). Bagwell and Ramey (1993) do not contemplate this possibility as in their model non-dissipative advertising cannot be informative if consumer and firms’ interests do not match: “Incentives for misrepresentation may exist [...], and, thus, costless quality claims will not always suffice to communicate quality.” Indeed, I show that costless (or cheap) quality claims *can be* informative even in the presence of incentives for misrepresentation.

Anand and Shachar (2007) provide a related model of advertising, where firms decide how much to invest in the precision of their messages which may or not be interpreted correctly by consumers. Although the present model does not focus directly on the investment in precision of advertising messages, it relates to that of Anand and Shachar (2007) in Section 2.5, where costly advertising is analyzed. In addition, the current model allows the firm

to choose the content of its communication policy freely, instead of having it constrained to either revealing its type or not.

Chakraborty and Harbaugh (2008) propose a related framework of cheap-talk under transparent motives. A sender is defined as having transparent motives if the receivers know which is the preferred action that the sender would like them to perform, but are uncertain about some variable that only the sender is informed about. In their paper Chakraborty and Harbaugh provide an application to advertising in a market for an ‘experience good’, in which consumers can only verify the product attributes after purchase. They show that when consumers are uncertain about both horizontal and vertical characteristics of the product, the firm can benefit from communication by emphasizing one of such dimensions (but not both). Unlike in their application, the current model focuses on a market for a ‘search good’ in which consumers gather more information about the product before buying. In addition, communication requires only one dimension of uncertainty: informativeness is attained because of the trade-offs that the firm has to face due to demand heterogeneity. Because no single advertising message can credibly attract all consumer segments in the market, the firm devises its advertising message in order to target its most profitable segment. Finally, in the present model the firm does not hold truly transparent motives: although consumers know that the firm wants to maximize profit, they do not know which is the target segment that the firm is trying to attract through its advertisement prior to credible communication taking place.

The rest of the paper is organized as follows. The next section presents the model and defines the equilibrium concept. Section 2.3 presents the main results. Section 2.4 analyzes the welfare effects of allowing misrepresentation, and Section 2.5 discusses how the results change when the cost of advertising is considered. Section 2.6 concludes.

2.2 Model

Consider the case of a monopolist firm endowed with a product of quality $q \in Q \equiv \{q_L, q_M, q_H\}$ where $q_H > q_M > q_L$. The product quality is its own private information and the probability of each of the quality levels taking place is equal to $1/3$. The inclusion of three quality levels allows us to consider informative semi-separating equilibria under pure messages.⁴

After the firm observes its product quality, it decides its price and its advertising strategy. Its message space coincides with Q , and its advertising message is observed by consumers prior search. There exist two segments of consumers (denoted by subscript $S \in \{L, H\}$) each of size one, who differ in their marginal valuations for quality $\{\theta_H, \theta_L\}$ such that $\theta_H > \theta_L \geq 0$, and in their outside options $\{t_H, t_L\}$, where $t_H > t_L \geq 0$. Consumers

⁴Note that a semi-separating equilibrium is also possible with only two types of firms when these are allowed to mix across messages.

observe the advertising strategy of the firm and decide whether to search or not by incurring a search cost $k > 0$. In case consumers decide to search, they get an individual specific gross utility draw $\theta_S q + \varepsilon_i$, where ε_i is assumed to be independent of the product quality, and follows the uniform distribution $U(-\gamma, \gamma)$ with $\gamma > 0$, and large enough such that the pricing solution is interior.⁵ In addition, search also enables consumers to observe the price of the product. Finally, consumers decide whether to purchase the product or not. In case they decide not to search or decide not to buy after the search cost has been incurred, they get some utility t_S from their outside option.

Two additional assumptions are required for informative communication to take place:

Assumption 1 (sufficient demand heterogeneity). *The marginal valuation for quality of the high type consumers θ_H is assumed to be strictly greater than twice the marginal valuation for quality of the low type consumers, θ_L .*

The power of this assumption ($\theta_H > 2\theta_L$) is further explained in Proposition 2. The intuition behind it is that firms need to be able to target different consumer segments through their advertising in order for communication to be credible. When there is not enough heterogeneity in demand, any advertising message has relatively the same effect on different consumer segments and informative communication is not possible. The constant ‘2’ is directly related to the mass of consumers of each segment. Intuitively, it is required to rise as the proportion of low type consumers increases, if credible communication is to occur.

Assumption 2 (low search costs). *Search costs are assumed not to be too large. In particular $k < \bar{k}$, where $\bar{k}$ is as defined in Section 2.3.*

The intuition behind requiring Assumption 2 is that when search costs are too high, no product is attractive in equilibrium whatever its quality, and advertising cannot be informative independently of the communication.

Finally, the boundaries of the quality range under analysis are defined such that the seller is willing to cater to all customer segments in case all of them search. A seller with a product of very high or very low quality would prefer to cater to a single customer segment: in case it sells a product of high quality, the seller may prefer to price the low type customers out of the market whereas if the product is of low quality the seller may prefer not to serve the high type customers since the quality of its product will be too low to cater to them at a reasonable price. The range for possible product quality $[q_L, q_H]$ is

⁵Details are described in the appendix.

defined such that all solutions are interior, and such that the firm is better off serving both types of consumers.⁶

Two comments about the basic setup of the model are in order. First, I assume prices are unobservable to consumers prior to search, and second, I limit the message space to the quality space Q . Both assumptions are also assumed by Bagwell and Ramey (1993) and are worth some discussion. First, it is often the case that final prices are not available to consumers before search. Even when they are, firms often offer extra conditions, warranties and add-ons that are unavailable to consumers until further search is realized. Other times there exist extra unadvertised fees that are only disclosed near the transaction phase. In addition, one can think of a similar model with observable prices where the price level is not sufficient to reveal all information about product quality. For example, in the case where consumers are uncertain about two or more attribute levels of a product (e.g.: performance and reliability), price may provide only partial information about the product and advertising may still have a useful signaling role for firms. Finally, even if price is observable, it will not be enough for separation if differences in product quality are not high enough.

Second, limiting the communication space to the quality space Q is a simplifying decision. Consider the case where price is also communicated. If the message $m = \{q, p\}$ is internally consistent, i.e., $p = \argmax \pi_q(p)$, then the quality component is a sufficient indicator for the price (consumers will have correct beliefs over prices given the communication). If on the other hand the message is not internally consistent (i.e.: p is not profit maximizing given the quality component q), then one would have to construct reasonable beliefs for such contradictory messages. One potentially reasonable belief would let consumers ignore such message (for example, a message of high quality and low price, or of low quality and high price). A simpler way of simplifying those beliefs would be to ignore one of the components (such as price), which is equivalent to the current approach. In the current model, rational consumers interpret message $m \in Q$ as a claim of a pair $\{q, p\}$, expecting to see high (low) quality claims associated with high (low) prices given the nature of the firm's profit maximization problem. The utility of consumer i belonging to segment S from buying is $u_{i \in S}(q) = \theta_S q + \varepsilon_i - p$. We can therefore think of the expression $\theta_S q + \varepsilon_i$ as the gross utility consumers get from the product. In fact, quality *does not ever need* to be observable after search as long as a draw from $f_{\theta q + \varepsilon}$ is realized for each consumer i , where one can think of quality as affecting the mean utility draw.⁷ In this case, more

⁶Expressions for the upper and lower bounds for quality are presented in the appendix. Additional cases can be analyzed outside of this interval but the relevant equilibria are equivalent to the ones found inside the range; yet the analysis is more complex.

⁷Nonetheless, in this case some utility draws provide consumers with enough information for them to invert for the true product quality.

quality is good news for all consumers in terms of gross utility, but it does not comprise all the information about the fit of the product with each customer.⁸

I consider the case where firms always prefer to communicate rather than not. This allows me to focus on the informational nature of the advertising strategy and can be justified when the benefits of advertising exceed its pure informational force (e.g. generating awareness, etc). I consider the advertising decision as sunk so as to isolate the informational effects. The analysis of the case where advertising implies a ‘small’ cost $\varepsilon > 0$ is presented in Section 2.5.

2.2.1 Equilibrium Concept and its Implications

The equilibrium concept I use is that of perfect Bayesian equilibrium.⁹ I focus on equilibria in which advertising is *informative* and *decisive*. Informative means that the consumers’ prior beliefs are updated by the existence of communication, whereas decisive means that the update in consumer beliefs due to communication makes search decisions different from the case where no communication occurs. An important remark is that an equilibrium where both types of consumers are willing to visit (at least) one firm-type cannot be both informative *and* decisive. Suppose there exists only one type of firm which both types of consumers would like to visit. In this case the remaining types of firms would prefer to deviate and claim to also be of that quality level, hence destroying informativeness (and decisiveness) in such an equilibrium. On the other hand when two out of the three firm-types can attract both consumer segments, the remaining firm would like to claim to be of one of the first types. In such case it is clear that in equilibrium misreporting cannot increase the expected consumer surplus of searching when compared to the case of absence of information. In this case the resulting equilibrium is not decisive. Therefore, no informative and decisive equilibrium can occur when all types of consumers visit the firm.

The first implication is that q cannot be in a given intermediate range located between q_L and q_H , in which both consumer segments would be willing to visit the firm. If such quality level q were possible, then all firms would like to claim to hold a product of that quality. This does not preclude a model of continuous quality, as long as we consider a continuum of consumers as well. The formal intuition of why there may not be a “perfect product” in the case of a continuum of product qualities is provided in the appendix.


The second implication of the equilibrium concept is that the set of candidate equilibrium prices is reduced to the ones that maximize profits when only one segment of consumers has visited, either L or H . Finally, throughout the paper I ignore ‘babbling

⁸Notice that $f(\theta_S q + \varepsilon_i | q)$ does not satisfy the strict monotone likelihood ratio property (SMLRP) under the current assumptions, and as such, although “more quality is good news”, it is so in a different sense from Milgrom (1981).

⁹In addition, only equilibria in pure strategies (where the firm sends a degenerate message) are analyzed, for the sake of clarity and of simplicity.

equilibria' whereby consumers ignore messages sent by the firm, leading the latter to be indifferent among messages.¹⁰ A summary of the timing of the model follows:

Figure 2.1: Timing



2.2.2 Purchase and Search Decisions

In this section I describe the stages of the model which are relevant for the purchase decisions of consumers, conditional on their search decisions. If the consumers of segment S decide to search, each consumer will then decide to buy if and only if $\theta_S q - p + \varepsilon_i > t_S$. Consider first the case where product quality q is known (or belief over q is correct) to customers before visiting, and price p is anticipated correctly (in which case the expected price p is a function of the belief over q). After searching, consumers decide whether to buy the product or not. For the ones who decide to buy, their expected surplus becomes

$$E_{\varepsilon} [\theta_S q - p + \varepsilon_i | \varepsilon_i > -(\theta_S q - p - t_S)] = \frac{1}{2} (\theta_S q - p + t_S + \gamma)$$

Consumers search if and only if the ex-ante consumer surplus unconditional on buying is higher than their outside option plus the search cost, which is equivalent to the condition:

$$\begin{aligned} E_{\varepsilon} (u_{i \in S} | q) &> k + t_S \\ \Leftrightarrow Pr(Buy|q)E(u_S|Buy, q) + Pr(Not Buy|q)E(u_S|Not Buy, q) &> k + t_S \end{aligned}$$

which becomes:

$$\frac{(\theta_S q - p + \gamma - t_S)^2}{4\gamma} > k \quad (2.1)$$

¹⁰I also ignore equilibria where consumers expect an arbitrarily high price which leads them not to visit. The absence of demand would ultimately lead the firm to be indifferent among any price level, including the price that was expected by consumers in the first place.

where the left hand side of the expression is the ex-ante expected consumer surplus net of the outside option t_S , and assumes the notation $CS_S(q)$ from here on.

When no information about quality exists, the participation constraint can be rewritten as follows:

$$\begin{aligned} E_{\varepsilon,q}(u_S) &> k + t_S \\ \Leftrightarrow \frac{1}{3}E_{\varepsilon}(u_S|q_H) + \frac{1}{3}E_{\varepsilon}(u_S|q_M) + \frac{1}{3}E_{\varepsilon}(u_S|q_L) &> k + t_S \end{aligned} \quad (2.2)$$

$$\Leftrightarrow \frac{1}{3}CS_S(q_H) + \frac{1}{3}CS_S(q_M) + \frac{1}{3}CS_S(q_L) > k \quad (2.3)$$

The ex-ante consumer surplus unconditional on quality level is additively separable into the consumer surpluses conditional on the possible quality levels. Accordingly, from the equilibrium beliefs and the Bayes rule one can write the appropriate expressions for the expected consumer surplus given a set of beliefs over product quality q .

2.2.3 Pricing Decision and Comparative Statics

Consider first the case where only one segment of consumers S visits the firm. The only candidate equilibrium price p^* maximizes the expected profit:

$$\pi_q(S) = p \frac{\theta_S q - p - t_S + \gamma}{2\gamma} \quad (2.4)$$

where $\pi_q(S)$ denotes the profit of the firm from serving solely consumer segment S . I focus on the case of interior prices which requires the support of ε_i to be sufficiently large.¹¹ The equilibrium price depends on the quality level of the firm and on which consumer segment decides to search. The only candidate equilibrium price consistent with customer segment S visiting is given by:

$$p^* = \frac{1}{2}(\theta_S q + \gamma - t_S) \quad (2.5)$$

The optimal price from equation (2.5) and the expressions for consumer behavior give rise to the following results (all proofs are in the appendix):

Proposition 1 (ex-ante consumer surplus comparative statics). *When a single customer segment S searches, its ex-ante equilibrium surplus of searching conditional on quality level q is increasing in both q and θ_S and decreasing in t_S .*

Proposition 1 makes intuitive sense: the product of the firm becomes more attractive as quality increases; in addition, as the outside option t_S increases, the search option loses value since both the probability of purchase as well as the expected surplus from purchase

¹¹See formal conditions in the appendix.

decrease. Finally, the expected surplus is not necessarily increasing with type S (since θ_S increases with S but so does t_S) which means that depending on the product quality, the low type consumers may be interested in visiting but the high types may not be, and vice-versa.

Let $CS_S(\mathcal{B}|p_S^*)$ denote the equilibrium ex-ante expected consumer surplus of customer $i \in S$, net of the outside option t_S , when she believes that $q \in \mathcal{B}$ and additionally that all the firm-types set prices targeting the segment S which are consistent with quality beliefs $\mathcal{B}$. Proposition 1 gives rise to the following corollary:

Corollary 1 (expected surplus ordering given beliefs). *When a single customer segment S searches, its ex-ante equilibrium surplus can be ordered in term of beliefs over q in the following order: $CS_S(q_L|p_S^*) < CS_S(\{q_L, q_M\}|p_S^*) < CS_S(\{q_L, q_M, q_H\}|p_S^*) < CS_S(\{q_M, q_H\}|p_S^*) < CS_S(q_H|p_S^*)$.*

Corollary 1 establishes that, in the case of a single customer segment, if consumers decide not to visit given a certain set of beliefs and expected optimal price, they will also not visit if beliefs are lower in terms of the ordering defined in the statement of the corollary.

Because equilibrium outcomes are defined by the absence of incentives to deviate by all agents, it is useful to describe how the surplus of consumers depend on product quality when they are not targeted by the store:

Proposition 2 (ex-ante consumer surplus comparative statics for non-participating consumers). *Given an equilibrium price targeting low type consumers, the ex-ante surplus of the high type consumers conditional on quality level q is always increasing in q . Additionally, under Assumption 1 ($\theta_H > 2\theta_L$), when the equilibrium price targets the high-type consumers, the ex-ante surplus of the low type consumers conditional on quality level q is decreasing in q for consumers of the low segment.*

The intuition for the second part of Proposition 2 is that the firm draws a substantial amount of profit from the high type consumers as quality increases. If the firm targets the high-type consumers through setting a high price, the benefits from quality increase slower than price for the low type consumers under Assumption 1, and ultimately leads them to avoid a product with too much quality since they will anticipate that its price will also be too high. Similar to Proposition 1, Proposition 2 gives rise to the following corollary:

Corollary 2 (expected surplus ordering given beliefs for non-participating consumers). *Under Assumption 1, the expected surplus can be ordered in term of beliefs over q in the following order: $CS_H(q_L|p_L^*) < CS_H(\{q_L, q_M\}|p_L^*) < CS_H(\{q_L, q_M, q_H\}|p_L^*) < CS_H(\{q_M, q_H\}|p_L^*) < CS_H(q_H|p_L^*)$ and $CS_L(q_L|p_H^*) > CS_L(\{q_L, q_M\}|p_H^*) > CS_L(\{q_L, q_M, q_H\}|p_H^*) > CS_L(\{q_M, q_H\}|p_H^*) > CS_L(q_H|p_H^*)$.*

As in Corollary 1, Corollary 2 establishes an ordering for consumer surpluses, this time for customers who do not visit the store in equilibrium.

2.2.4 Demand Heterogeneity and Consumer Surplus

This section describes how consumer heterogeneity affects the surplus of consumers themselves. Which consumers benefit from the presence of other consumers? Although there are no informative and decisive equilibria when both types of customer segments visit, a result from such Bayesian equilibrium is presented in this section. Let $\tilde{q}$ be the quality threshold that satisfies the equation $\pi_{\tilde{q}}(L) = \pi_{\tilde{q}}(H)$, where π_q is as defined in equation (2.4). At quality level $\tilde{q}$ the firm is indifferent between exclusively serving either one of the customer segments. The expression for such quality level, given that an interior price exists, is:

$$\tilde{q} = \frac{t_H - t_L}{\theta_H - \theta_L} \quad (2.6)$$

Expression (2.6) makes intuitive sense: at quality level $\tilde{q}$ the additional utility that high type consumers get from the product over what the low types get, $(\theta_H - \theta_L)\tilde{q}$, is equal to the difference in opportunity costs, $t_H - t_L$. At this quality level both customer segments are equally attractive for the firm. If it had to pick only one segment to serve, a firm with quality $q < \tilde{q}$ would prefer the low type customers whereas a firm with quality $q > \tilde{q}$ would prefer the high type customers. The introduction of the following definition conceptually helps in the subsequent analysis:


Definition 1 (target customer segment). Segment S is defined as the target segment of a firm with product quality q whenever $\pi_q(S) > \pi_q(\sim S)$, $S \in \{L, H\}$.

The following proposition follows:

Proposition 3 (simultaneous demand competition and complementarity). *The target segment of a firm with product quality q prefers the equilibrium where the non-target segment also visits the firm. In addition, the non-target segment prefers the equilibrium where the target segment does not visit the firm.*

Proposition 3 holds the following interpretation: In the first case where $q < \tilde{q}$, the target customers L prefer the outcome with consumer heterogeneity whereas the non-target customer segment H prefers to be the exclusive customer segment. Similarly, when $q > \tilde{q}$ the target customer segment H prefers the outcome with consumer heterogeneity whereas the non-target customers of segment L prefer exclusivity. This is a somewhat surprising yet intuitive result upon reflection: the most profitable customers prefer the outcome where they are not the exclusive segment because they know that the lack of consumer heterogeneity will let the firm extract higher rents from them. On the other hand, the least profitable customer segment would like to be served alone: otherwise, whenever the target customers also ‘show up in the store’, equilibrium dictates that the firm will have its price mainly extracting rents from the most profitable customer segment and the surplus of the non-target consumers will decrease at a high rate. In both scenarios there exists simultaneous “consumer competition” and “consumer complementarity” for different customer segments. Figure 2.2 summarizes the findings presented until this point. The consumer surplus for the high type consumers is always increasing in quality. However, the consumer surplus for low-type consumers may be decreasing in quality if the firm has set price level p_H^* . Notice that when $q < \tilde{q}$, both consumer segments prefer to face price p_H^* rather than p_L^* . This is because when quality is low, $p_L^* > p_H^*$ since the firm would have to set a very low price (p_H^*) in order to profitably sell a low quality product to the high type consumers. When quality is high ($q > \tilde{q}$) the firm’s product is relatively more attractive to high type consumers, leading to $p_H^* > p_L^*$.

Figure 2.2: Ex-ante consumer surplus




$CS_S(p)$ denotes the ex-ante expected consumer surplus of customer segment S from searching, net of the outside option, when facing price p . Parameter values are $\theta_H = 5/8$, $\theta_L = 1/4$, $t_H = 1/10$, $t_L = 11/200$, $\gamma = 11/100$, $k = 3/1000$.

All consumer surplus curves intersect at quality threshold $\tilde{q}$, independently of the consumer segment and the equilibrium price. This is due to the nature of $\tilde{q}$, which equates the benefits of quality and opportunity costs across consumer segments.

2.3 Advertising Strategies and Market Outcome

As discussed in section 2.2.1, there cannot be informative and decisive equilibrium for an intermediate region of q .

Figure 2.3: ex-ante expected consumer surplus when facing price p .



$CS_S(p)$ denotes the ex-ante expected consumer surplus of customer segment S searching when facing price p . Parameter values are $\theta_H = 5/8$, $\theta_L = 1/4$, $t_H = 1/10$, $t_L = 11/200$, $\gamma = 11/100$, $k = 3/1000$.

The quality thresholds $q_{(1)}$ to $q_{(3)}$ in Figure 2.3 indicate some useful intersections between the ex-ante consumer surplus curves and the search cost k , and are defined in the appendix. The intermediate quality range where there cannot be informative and decisive equilibria is given by $[q_{(2)}, q_{(3)}]$. If there was a strictly positive probability of a firm having product quality $q \in [q_{(2)}, q_{(3)}]$, then all firms would like to claim to hold such a product, and communication would break down.

Secondly, there can be no informative and decisive equilibria where all firms are below or above $\tilde{q}$. This is because in this case all the firms are interested in the same consumer segment, and decisiveness cannot be attained.

Finally, for effective communication to take place it must also be that the search cost k is sufficiently low; otherwise advertising cannot provide decisive information to consumers. The upper-bound for the search cost $\bar{k}$ is defined as $CS_S(\tilde{q}|p_{S'}^*)$ where $S, S' \in \{L, H\}$. The intuition behind this is that any firm with a product quality $q < \tilde{q}$ is not worth visiting for

any consumer in equilibrium. In that sense all the firms will try to behave as if they were above $\tilde{q}$ (and targeting the same consumer segment), and decisiveness breaks down.

2.3.1 Matching Interests

Let $CS_S^*(\mathcal{B})$ be the ex-ante expected surplus of consumers of segment S , net of the outside option t_S , when they believe that $q \in \mathcal{B}$. In this case, consumers of segment S search whenever $CS_S^*(\mathcal{B}) > k$. Consider the conditions for the candidate equilibrium of truthful communication, referred to as (C1):

- $q_L, q_M < q_{(2)}$ and $q_H > q_{(3)}$.
- Firms q_L, q_M set price p_L^* while firm q_H sets price p_H^* .
- $CS_L^*(q_L) > k, CS_H^*(q_H) > k$.
- $CS_L^*({q_L, q_M, q_H}) < k$ or $CS_H^*({q_L, q_M, q_H}) < k$ (so that the equilibrium is decisive for at least one consumer segment).
- Beliefs are: $q_L^m \rightarrow \hat{q}_L; q_M^m \rightarrow \hat{q}_M; q_H^m \rightarrow \hat{q}_H$ (i.e.: consumers believe the messages of the firm).

Under the profile above, each type of firm sets prices in order to target its preferred consumer segment. Low type consumers are willing to visit a low quality firm (and hence also willing to visit a medium quality firm), whereas high type consumers are willing to visit a high quality firm. Finally, at least one of the consumer segments is not willing to visit under absence of information, and consumers believe the claims of the firm to be true. Under the conditions above we get the following outcome:

Proposition 4 (equilibrium with matching interests) *Under (C1), there exists a truthful communication equilibrium where each type of firm advertises its type, and only low type consumers search whenever the firm is of low or medium quality, whereas only high type consumers search whenever the firm is of high quality.*

Under the conditions above there is no incentive for firms to misreport their true product quality. In that case, each type of firm targets its preferred consumer segment: firms q_L and q_M target the low type consumers - which is their most profitable segment since their quality is low, i.e.: $q_L, q_M < \tilde{q}$ - while firm q_H targets the high type consumers - since it is its most profitable segment by $q_H > \tilde{q}$. Credibility arises because of the match between the interests of firms and consumers. This is because the low and medium quality firms hold a product that is unattractive to high type consumers, while the high quality firm holds a product that is too expensive for low type consumers. The fact that under absence of communication at least one consumer segment is unwilling to search is changed by advertising.

Other equilibria with matching interests can occur. First, because the firms of types q_L and q_M are indifferent between being believed to be of low or medium quality (since they can both attract the low type consumers), another possible outcome is that they both use the same advertising message. Consumers will adjust their beliefs accordingly and will visit the firms (the details are provided in the appendix). Second, in the case where $q_L < q_{(2)}$ and $q_M, q_H > q_{(3)}$, we get the symmetric equilibrium to the one described in Proposition 4, where firms also report their true product qualities truthfully.

Both equilibria with matching interests hold the same intuition: the lower type firms would rather advertise truthfully instead of overstating their quality levels, since otherwise they would attract a ‘demanding’ high segment that has access to higher outside options and becomes ‘disappointed’ after inspecting the product. In addition, such advertising would ‘scare away’ the low type consumers who would anticipate too high of a price were they to believe the advertising message. On the other hand, the higher type firms prefer to attract the high type customers since these have a higher willingness-to-pay for quality. For those firms, understating product quality could only lead to attracting low type customers who would not be willing to pay as much for quality, while it would simultaneously scare the high-type customers away.

2.3.2 Incentive to Misreport

As in Crawford and Sobel (1982), there exists an informative (and decisive) equilibrium despite the existence of incentives to misreport by the firm. This occurs as long as not all types of firms would like to be recognized as being of the highest quality (this is equivalent to the sender bias ‘ b ’ being small in the example provided by Crawford and Sobel). That case can occur under the following conditions (referred to as C2):

- $q_L, q_M < q_{(2)}$ and $q_H > q_{(3)}$.
- Firms q_L, q_M set price p_L^* while firm q_H sets price p_H^* .
- $CS_L^*(q_L) < k$, $CS_L^*(q_L, q_M) > k$, $CS_H^*(q_H) > k$.
- $CS_L^*({q_L, q_M, q_H}) < k$ or $CS_H^*({q_L, q_M, q_H}) < k$ (so that the equilibrium is decisive for at least one consumer segment).
- Beliefs are: $q_L^m \rightarrow \widehat{q_L}$; $q_M^m \rightarrow (\widehat{q_L}, 1/2; \widehat{q_M}, 1/2)$; $q_H^m \rightarrow \widehat{q_H}$.

In this case the low type consumers are not willing to search if they believe that the firm is of low quality, but are willing to ‘take the chance’ when they cannot distinguish whether product quality is low or medium. In this case we get the following outcome:

Proposition 5 (equilibrium with incentive to misreport) *Under (C2), there exists an equilibrium with incentive to misreport, where the low quality firm claims to be of medium*

quality, the remaining firms report truthfully, only low type consumers visit whenever they see medium quality being advertised, and only high type consumers visit when high quality is advertised.

In this case, since only the high quality firm has an incentive to report q_H , the consumers know that any other message can only come from q_L or q_M . Firm q_L strictly benefits from exaggerating its quality level by stating q_M in order to attract the low type consumers. This strategy will work because firm q_M has an incentive to report its quality truthfully, and so consumers cannot distinguish the lower types of firms through their advertising. The reason why such a mechanism can induce search is because such a message also provides information: consumers expect a firm of low or medium quality when they observe message q_M , and additionally expect the product not to be of high quality, and hence not to be too expensive. This last piece of information is essential for the equilibrium to be decisive: otherwise, the effect of overstating product quality could not induce consumers who would otherwise be unwilling to search to do so.

There are no other equilibria with mismatching interests. For example, it is never the case that the firm with medium quality prefers to misreport and pretend to be of type q_H whereas firm q_L prefers to report truthfully. This case would require that $q_L < \tilde{q}$ and $q_M, q_H > \tilde{q}$. But when $k < \bar{k}$, any quality above $\tilde{q}$ is satisfactory for the high type consumers and so no mismatch of interests occurs.

Finally, it is also never optimal for a firm to understate its true quality level. To understand why consider such a case. The firm will only do this if it wants to target the low type consumers (since only these may benefit from a lower quality product), in which case it must be that $q_M < \tilde{q}$. In this region however, less quality is always bad news for low type consumers and so it follows that such firm cannot strictly benefit from underreporting its true quality level.

2.4 Welfare Implications

This section analyzes the impact of allowing for misrepresentation vis-a-vis the case of perfect information. Table 2.1 presents ex-ante surplus of the low consumer segment and the profit of the low-type firm (the surplus and profits of the remaining agents is unchanged by the advertising policy).

Table 2.1: Ex-ante consumer surplus for different advertising policies

	Perfect information/Truth telling	Misrepresentation
$E[CS_L]$	$\frac{2}{3}t_L + \frac{1}{3} \left[\frac{(\theta_L q_M - t_L + \gamma)^2}{16\gamma} - k \right]$	$\frac{2}{3} \left[\frac{1}{2} \frac{(\theta_L q_L - t_L + \gamma)^2}{16\gamma} + \frac{1}{2} \frac{(\theta_L q_M - t_L + \gamma)^2}{16\gamma} - k \right] + \frac{1}{3}t_L$
π_{q_L}	0	$\frac{(\theta_L q_L - t_L + \gamma)^2}{8\gamma}$

Whenever an incentive for misrepresentation exists, consumers never benefit from the firm being able to exaggerate its product quality. The reason why the firm profits from misrepresentation is because consumers are willing to visit in the hope of finding a good fit given that the firm may be of low or medium product quality. On the other hand, the firm benefits from misrepresentation since after being able to induce visits, the search costs are sunk, and a significant proportion of consumers may decide to buy. It follows that whenever the increase in profits by the firm exceeds the loss in consumer surplus by consumers, total welfare may actually increase when firms are allowed to misrepresent themselves.

2.5 Extensions

2.5.1 Costly Advertising

In this section I consider a small cost of advertising $\varepsilon > 0$, which firms have to incur in order reach consumers with their message. This cost can affect firms' strategies even if it is assumed to be as small as needed so that it does not affect the communication decision by its scale, but only by its existence. The effective message space becomes $Q \cup \{\emptyset\}$, where message $\emptyset$ is defined as the non-informative option. This assumption allows us to also analyze the discrete advertising decision as well as the content of communication.

2.5.1.1 Matching Interests

Consider the same equilibrium as in the case with matching interests presented in Proposition 4. Two equilibria are possible. The first is supported by the belief that no message ($\emptyset$) means "high type". In this case the high type firm prefers not to advertise whereas the low and middle type firms advertise their own types by incurring the advertising cost ε (as before, they are indifferent between messages q_L and q_M). If consumers believe that no communication means quality levels q_L or q_M on the other hand, then only firm q_H advertises its type in order to attract the high type consumers.

The intuition behind the effects of the introduction of an advertising cost is simple: whenever 'not advertising' conveys a belief that is the most advantageous for a firm, such a firm will prefer not to incur the advertising cost. That possibility appears because other firms who do not share the same benefit have to incur advertising costs, letting the first firm

‘free-ride’. Which beliefs are applicable to the no-communication case cannot be addressed here. Only a prediction towards asymmetry of advertising strategies is provided.

2.5.1.2 Incentive to Misreport

The analogous equilibrium to that of incentive to misreport (Proposition 5) also holds under costly advertising. As before, an asymmetric advertising case occurs when the absence of communication is synonym of high quality. In this case the high type firm does not advertise while firms q_L and q_M both want to advertise medium quality in order to attract the low type consumers. On the other hand, when no communication is interpreted as ‘the firm is not of high type’, then only the high-type firm communicates and there is no need for the remaining firms to incur advertising costs. Finally, when no communication is interpreted as the firm being of low type, then all firms are better off advertising. Firms q_L, q_M will advertise middle quality while firm q_H communicates its quality level correctly. Notice that only in the case of incentive to misreport is there an equilibrium where all firms advertise. This happens because the firm who does not communicate is interpreted to be of an unattractive type for all consumers.

2.6 Conclusion

The goal of this paper is to describe how advertising can be informative in the context of a vertically differentiated market. The fact that different types of firms may have different preferences over consumer segments is central for credibility, and no ‘money burning’ is required for informative communication to take place. This is consistent with some empirical findings. For example, Caves and Greene (1996) find no relation between advertising expenditures and product quality after considering a sample of 200 products analyzed by Consumer Reports.

In the present mechanism, a firm with a low quality product may not want to exaggerate its claims by too much since it may scare some of its potential customers away because of high price expectations, while it will attract consumers who will be disappointed upon search. Similarly, a firm with a high quality product may not want to understate its true quality level since it will attract consumers who are willing to pay too little for quality while discouraging others who would be willing to pay more.

The current model recognizes that even in the context of vertically-differentiated markets there exists a matching problem between firm and consumers. Advertising can contribute to solve the matching problem. When the interests of firms and consumers are aligned, truthful communication takes place in the market. In addition, the model describes how the existence of incentives to misreport do not necessarily make communication break down as long as they are not too large. In this case a firm can benefit from misreporting its

true quality level and induce consumers to search. This leads differentiation to be sparser in the lower partition than in the higher one.

Effects of allowing misrepresentation in advertising policies on welfare vary. Although misrepresentation always hurts consumers, it may benefit the firm to the extent that the net effects will be positive overall. When the low quality product is much worse than the medium quality product, allowing for misrepresentation in advertising hurts consumers more than it generates profits for the firm. On the other hand, when low and medium product qualities are not too far apart (and hence exaggeration not too big), total welfare may actually increase through misrepresentation. From a policy perspective, exaggeration should thus be allowed, but only in “fair amounts”.

While the analysis above considers the case of a monopoly, the setting with competition may also yield interesting results. In addition, considering the existence of several product attributes, the possibility of a continuum of quality levels and the analysis of the optimal advertising strategy along the product life-cycle may prove to be worthwhile challenges for future research.

Chapter 3

Information as a Homogenizing Force: The Impact of Others' Opinions in Online Ratings¹

3.1 Introduction

With the rise of the Internet, information has gained a new dimension: it is often provided and used by several agents. While the previous chapters focused on information in a setting of transmitter-receiver, this chapter looks at a setting where the information users are also information providers to other subsequent agents.

The Internet has become a popular source of information gathering as well as information sharing. Most websites such as *amazon.com*, *netflix.com* and *target.com*, that sell or rent products, also let consumers share their consumption experience. Even when the retailers do not offer rating services, websites such as *imdb.com* and *yelp.com* let consumers share their consumption experiences about virtually any product and service. Consumers routinely check the information available on these websites before making a purchase decision for a product with uncertain quality regardless of whether the purchase is made online or through other channels. The same consumers often return to their favorite websites to share their product consumption experiences.

Online ratings and reviews have become a popular way of reporting a product consumption experience on the Internet. The following observations give us an idea of how common the use of online ratings is: in the first 4 weeks, '*Inglourious Basterds*' received over 60,000 ratings on *imdb.com*; '*Harry Potter and the Deathly Hallows*' has received 3,410 reviews on *amazon.com*; and '*The Venetian Las Vegas*' has received 1,321 reviews on *priceline.com*.²

¹This chapter is based on co-authored work with Shubhranshu Singh.

²As of 10/01/2009.

When reporting these ratings and/or reviews consumers try to incorporate all the information about the quality of the product that they have acquired from various sources including, but not limited to, their own product consumption experience.

Websites vary in how they let consumers report their consumption experience. Some websites such as *yelp.com* and *target.com* require users to write reviews together with ratings; others such as *imdb.com* and *walmart.com* let consumers decide whether they want to write a review or just rate the product on an ordinal scale. One major difference between reviews and ratings, apart from the fact that ratings take far less time to report, is that while ratings are mostly anonymous reviews are not. This might have implications on how consumers behave when entering online ratings compared to writing online reviews. We focus our attention on the online ratings in this paper.

Extensive research has been carried out in the study of information transmission. According to this research, information transmitters report extreme attitudes (Cohen, 1961), look for supportive information and ignore inconsistencies (Brock and Fromkin, 1968). Transmitters are also shown to modify their communication in the presence of self-presentation concerns. Schlosser (2005) demonstrated through experiments that "posters" (those communicating their experiences to others) appearing later in a sequence of communications get influenced by earlier negative opinions (but not positive opinions) and adjust their public attitude downwards when responding in public because of self-presentation concerns. Those who reported their opinions privately were less influenced by negative opinions as they had no self-presentation concerns.

It is common for consumers to get exposed to the average of the prior ratings, prominently displayed on screen, immediately before reporting their own rating. The literature on information transmission may be more inclined towards the possibility that consumers will ignore the average rating. Based on Cohen (1961), we expect consumers to report extreme ratings, which makes it less likely that their ratings would be affected by the average of the prior ratings. Consumers might find the average of the prior ratings to be inconsistent with their opinion in some cases and consistent with their opinion in other cases. Following Brock and Fromkin (1968) we expect consumers to disregard the average ratings that are inconsistent with their opinion. Finally, since there are little self-presentation concerns in online ratings (because they are usually anonymous), the results by Schlosser (2005) suggest that those arriving later in the sequence would not change their ratings upon observing the average of the prior ratings.

There are several ways in which consumers could use the average of the prior ratings to arrive at their final ratings. One possibility is that consumers might have an objective of guiding the audience to the same opinion as their own. This would require them to enter extreme low ratings if the average of the prior ratings is higher than their opinion, and vice-versa.

Another possibility is that in order to arrive at the final ratings consumers might start with the displayed average ratings and then adjust until they reach a value they think is acceptable.³ Tversky and Kahneman (1974) refer to this mechanism as the *anchoring and adjustment* heuristic. Further research (e.g. Plous, 1989; Jacowitz and Kahneman, 1995; Mussweiler and Strack, 2004) has shown that final estimates provided by subjects are influenced by initial anchor values. For example, subjects estimated that Gandhi lived to be 67 years old when they first answered whether he died before or after the age of 140, but they estimated that he lived to be only 50 years old when they first answered whether he died before or after the age of 9 (Strack and Mussweiler, 1997). Although anchoring and adjustment is helpful in most cases, adjustments tend to be insufficient, leaving people's final estimates biased towards the starting values (see also Epley and Gilovich, 2006).

Yet another possibility is that consumers might consider the average of the prior ratings as relevant information and intentionally incorporate it in their evaluations in the process of arriving at their final ratings. Individual consumers might not consider themselves experts in evaluating all the aspects of the products. This might prompt them to incorporate the information present in the average of the prior ratings in their own evaluations in the hope of providing better evaluations of that product. Some examples in which a decision maker's action is influenced by those of previous decision makers follow. People choose to dine in restaurants where they see a higher fraction of seats occupied (Becker, 1991); voters are known to be influenced by opinion polls to vote for the party that is expected to win (Cukierman, 1991); '*The Discipline of Market Leaders*' remained selling as best seller for some time once authors secretly bought 50,000 copies of the book to make the *New York Times* bestseller list (see Bikhchandani et al., 1998); researchers are known to work on topics that are hot (see Banerjee, 1992); Patients adjust their quality perception about a kidney for transplant based on refusals by patients ahead in the waiting line (Zhang, 2010). By use of movie ratings data from *netflix.com*, we investigate whether exposure to the average of the prior ratings displayed on screen leads individuals, intentionally or unintentionally, to incorporate it in their own ratings. We are also interested in the impact of such behavior on the overall pattern of the ratings. Ariely et al. (2003) reported in their working paper that subjects were willing to pay \$11.62 for a wine with a *wine advocate* rating of 86 compared to \$17.42 for the wine with a *wine advocate* rating of 98. This suggests that consumers' willingness to pay should be higher for a product with a higher average rating. To the extent that this result is generalizable, firms may try to take advantage by including inflated ratings. The fact that *imdb.com* has adopted a weighted average scheme to offset such efforts highlights the importance of this issue. ⁴We therefore also examine to which extent it is possible to artificially inflate the average rating of a movie in our dataset.

³This is especially true in the presence of uncertainty.

⁴From *imdb.com* (10/08/09) on the weighted average rating scheme: "The objective of the scheme is to present a more representative rating which is immune from abuse by subsets of individuals who have combined together with the aim of influencing (either up or down) the ratings of specific movies. This

We use an ordered logit model to test our hypothesis since our dependent variable, ratings, is on an ordinal scale. We treat the average of the prior ratings and movie dummies as the independent variables and let the unobserved factors follow a logistic distribution. The rating given by an individual might also depend on the media reports, on the release of competing titles and on other marketing efforts by movie studios, which are not observed. These shocks might have a carryover effect on consumers who rate the movie in subsequent periods. Not accounting for this carryover effect could bias the coefficient. We therefore add a movie specific error term and allow it to be serially correlated. We note that estimation of limited dependent variable models with autocorrelated errors has received relatively little attention probably due to computation complexity of obtaining the maximum likelihood estimator. We use a Bayesian approach in order to implement a feasible estimation method.

In terms of methodology, this paper draws heavily from Albert and Chib (1993) and Allenby and Lenk (1994) in order to estimate an ordered logit model with autocorrelated errors. Albert and Chib (1993) propose a Bayesian method for modeling categorical response data using an ordered probit model. Applications of this approach in marketing include Manchanda and Gupta (1999), Jedidi and Manchanda (2003) and Chung and Rao (2003). Allenby and Lenk (1994) apply a logistic normal regression with autocorrelated errors using Bayesian procedures to scanner panel data of ketchup purchases. McCulloch and Rossi (1994) also discuss in short a binary multiperiod probit model with autoregressive errors.

We find that the coefficient of the average rating is positive and significant. This implies that when reporting anonymous online ratings, individuals incorporate the average ratings displayed on the screen in their evaluations of the movie. Our findings are in contrast with those of Schlosser (2005), who reports that in an anonymous setting, feedback is not affected by previous responses. One possible explanation could be that in Schlosser's setting, subjects were experts and they might not see the value of incorporating others' feedback in their own rating. However, a typical consumer might not consider herself an expert and may treat the average of the previous ratings as a relevant piece of information. Simulations indicate that incorporating the average of the prior ratings in one's own evaluations changes the overall pattern of ratings for a representative movie. We also find that a few early extreme ratings, good or bad, change the ratings that a movie receives in subsequent periods. This effect creates a lasting impact on the average rating of the movie. Therefore, the final ratings of a movie may critically depend on a handful of initial ratings. The rest of the paper is organized as follows. Section 2 summarizes the data. Section 3 describes the ordered logit model with serially correlated errors. Section 4 presents the insights into rating behavior of individuals from the model and from the estimation procedure. Section 5 summarizes the findings.

includes people involved in the production of a movie and their friends or fans trying to unduly raise the rating of a movie far above that of where the typical IMDb users would rate it."

3.2 Data

The use of online ratings is widespread. Due to the availability of the data we chose online movie ratings for this study. People typically refer to one or more movie websites like *rottentomatoes.com*, *imdb.com* or *movies.yahoo.com* to look at movie ratings, among other things, before making a decision to watch movies. Many of them return to rate the movies they watch. Most movie review websites let their registered users rate movies. While they use an ordinal scale, the number of possible ratings varies from website to website. *imdb.com* allows for 10 possible ratings from 1 to 10; *movies.yahoo.com* allows for 13 possible ratings from F to A+; and *rottentomatoes.com* allows for 11 possible ratings from 0% to 100%.

For this study we use the *netflix.com* (an online DVD rental website) dataset.⁵ Movies received an average of over 5,000 ratings during this period. For this study, a sample of 60 movies was randomly chosen from a subset of movies that received on average at least one rating per day. These 60 movies were released between 1934 and 2004 and represent all genres.

We also note that *netflix.com* users are often exposed to the average of the prior ratings but not necessarily to the number of ratings. Although the latter information is available, the user has to make an effort if she wants to know the number of ratings based on which the average is reported.

Cumulative average ratings were calculated for each movie for each period using the individual ratings data. We round the average ratings to the first decimal in order to be consistent with how they appeared on screen. Average ratings are as low as 1.9 and as high as 4.8 in the working sample. Only 200 periods of data are included in the study as the average ratings typically do not change after 200 periods. Figure 3.1 presents average ratings over 200 periods for three movies in the sample. It can be seen that while average ratings changed frequently in early periods, they become stable after about 150 periods. We therefore restrict the data to only 200 periods.

⁵Data description by Netflix: "The movie rating files contain over 100 million ratings from 480 thousand randomly-chosen, anonymous Netflix customers over 17 thousand movie titles. The data were collected between October, 1998 and December, 2005 and reflect the distribution of all ratings received during this period. The ratings are on a scale from 1 to 5 (integral) stars."




Figure 3.1: Average ratings changing over time for three randomly chosen movies.




Figure 3.2: Distribution of ratings for 5 randomly chosen movies.

It has been observed that most individuals choose to report their ratings either when they are highly satisfied or when they are highly dissatisfied with the product. This might create bias in our results. However, we did not observe this bimodal distribution in Netflix data. Distributions of ratings for five randomly chosen movies are shown in Figure 3.2. One of the reasons for this fact could be that since Netflix users get movie recommendations based on the ratings they give to other movies, they have an incentive to report all experiences; not just the extreme ones. This may offset the urge to influence others' opinions through the use of extreme ratings.

3.3 The Model

Suppose that all the individuals rating the movie m at time t receive the same signal. This is a simplifying assumption and may be relaxed in a future version of the paper. We transform all the ratings given to movie m in period t to one rating on the same ordinal scale of 1 to 5, by taking the average and rounding it in the following fashion. We assume that if the average of the ratings given on a particular day for a movie is 2.3 for example, a representative consumer would give a rating of 2 with probability 0.7 and a rating of 3 with probability 0.3. ⁶Suppose $r_{m,t}$ denotes the rating for movie m at period t given by a representative consumer. This rating could take one of the five possible values (from 1 to 5) represented by $j \in \{1, 2, \dots, J\}$. Let $y_{m,t}^*$ be the latent utility for movie $m \in \{1, 2, \dots, M\}$ at rating occasion $t \in \{1, 2, \dots, T_m\}$.

There exist cutoff points $\{k_0, k_1, \dots, k_J\}$ which divide the space of $y_{m,t}^*$, the real line, in J intervals, satisfying

$$-\infty < k_0 \leq k_1 \leq \dots \leq k_J \equiv +\infty$$

such that

$$r_{m,t} = j \text{ if } k_{j-1} \leq y_{m,t}^* \leq k_j.$$

We decompose the latent utility as

$$y_{m,t}^* = \delta f(R_{m,t}) + (1 - \delta)(\alpha_m d_m + \epsilon_{m,t} + \xi_{m,t}) \quad (3.1)$$

where δ is the weight that individuals put on the average of the prior ratings, $R_{m,t}$ represents the average of the prior ratings for movie m at period t and d_m represents the movie dummy. The error term $\xi_{m,t}$ is distributed logistic with cumulative distribution

⁶We also used a more general method of drawing the ratings by using the daily empirical distribution of the ratings and obtained similar parameter estimates.

$$F(\xi) = \frac{\exp(\xi)}{1 + \exp(\xi)}$$

and $\varepsilon_{m,t}$ is an error term, potentially serially correlated. These errors allow the model to accommodate systematic changes in opinions of movies. For example, the public opinion for a movie might be favorable at a specific moment in time, leading the average rating to increase. Ignoring that these opinions may be correlated over time could lead us to attribute a disproportionate influence of the average rating to the next period's ratings.

The errors $\{\varepsilon_{m,t}\}$ are independent across different movies. The error terms within movie m are autocorrelated as:

$$\varepsilon_{m,t} = \phi_m \varepsilon_{m,t-1} + \zeta_{m,t}$$

for $t = 1, 2, \dots, T_m$.

The factor ϕ_m represents the autocorrelation parameter, and is between -1 and 1. $\zeta_{m,t}$ is independent and identically distributed $N(0, \sigma^2)$. The error term in period $t = 0$, $\varepsilon_{m,0}$, is independent and identically distributed $N(0, 1)$.

The function $f(\cdot)$ transforms the average of the prior ratings to utilities. $f(R_{m,t})$ is given by the inverse logistic cumulative distribution

$$f(R_{m,t}) = s_m \log \left(\frac{R_{m,t} - 1}{5 - R_{m,t}} \right)$$

where s_m is the scale parameter of the logistic distribution that is used to transform the average of the prior ratings to the utilities. The parameter s_m is restricted such that the variance of this logistic distribution is same as the variance of the latent utility in equation (3.1).

The probability of rating j being given to movie m at time t simplifies to

$$p_{m,t}(j) = \frac{\exp\left(\frac{k_j}{1-\delta}\right) - y_{m,t}}{1 + \exp\left(\frac{k_j}{1-\delta}\right) - y_{m,t}} - \frac{\exp\left(\frac{k_{j-1}}{1-\delta}\right) - y_{m,t}}{1 + \exp\left(\frac{k_{j-1}}{1-\delta}\right) - y_{m,t}}$$

for $j = 1, 2, \dots, J$; $t = 1, 2, \dots, T_m$ and $m = 1, 2, \dots, M$, where $y_{m,t}$ is given by

$$y_{m,t} = \frac{\delta}{1-\delta} f(R_{m,t}) + \alpha_m d_m + \varepsilon_{m,t}.$$

Hereafter, we refer to the coefficient $\frac{\delta}{1-\delta}$ as α_1 and to the vector formed by α_1 and α_m as α .

3.3.1 The Gibbs Sampler

We estimate the autoregressive ordered logit model by generating random draws from the model's posterior distribution. This distribution does not have a convenient form from which we can take random draws. We use the Gibbs sampling, a Markov Chain Monte Carlo (MCMC) method for taking draws from the joint posterior. Gibbs sampling proceeds by drawing iteratively from the conditional posterior densities. As the number of iterations approach infinity, the process converges to draws from the joint posterior distribution.

Using the convenient notation $[A|B]$ for representing the conditional distribution of A given B introduced by Gelfand and Smith (1990), we can write the joint distribution of data and parameters as

$$[R, Y, \varepsilon_{\cdot,t}, \alpha, \Phi, \sigma^2, K] \propto [R|K, Y][Y|\varepsilon_{\cdot,t}, \alpha, \Phi, \sigma^2][\varepsilon_{\cdot,t}][\alpha][\Phi][\sigma^2][K]$$

where $r_{m,t}$ and $\{y_{m,t}\}$ are stacked to obtain R and Y . All observations for a movie appear in chronological order followed by all observations for the next movie. $\varepsilon_{\cdot,t}$ is a vector of $\varepsilon_{m,0}$ and Φ is a vector of ϕ_m in the same order as movies appear in R and Y . K is a vector of cutoff parameters $\{k_0, k_1, \dots, k_J\}$ where $k_0 \equiv -\infty$ and $k_J \equiv \infty$. To ensure that parameters are identifiable, it is necessary to impose one restriction on the cutoff parameters; without loss of generality we assume $k_2 \equiv -0.5$. The posterior conditional distribution from which random draws are taken is described in detail in the appendix. They are presented in the same order in which draws were taken in each iteration.

3.3.2 Simulation


To test our algorithm, we consider a simulated example. We also want to look at the time it takes to run an iteration and if the sampler is sensitive to initial conditions. We generated the data assuming a set of parameter values, $(\{k_0, k_1, k_2, k_3\}, \alpha_1, \phi, \sigma^2) = (\{-\infty, -1, 1, \infty\}, 1, 0.5, 1)$. (For simplicity we assume the rating can take only three values instead of five.)

As shown in Figure 3.3, the Gibbs sampler converged quickly. The convergence plots are shown for the Gibbs sampler with starting values = (0.5, 0.5, 0, 0.5). The parameter was fixed at -1 for identification. The parameter values recovered from the sampler were = (1.03, 0.91, 0.53, 0.95) and standard errors were (0.02, 0.03, 0.18, 0.07) respectively. We verified the sensitivity of the sampler using different sets of starting values and got very similar results to the one show in Figure 3.

3.4 Estimation Results

We start with the basic model (Model 1) with fixed effects and independent error terms and allow for serially correlated errors terms in our full model (Model 2). Both the models were estimated with 75,000 iterations each. 50,000 iterations were used for burn-in and the last 25,000 iterations were used to calculate the mean and standard error for each run.

Figure 3.3: Convergence plots using simulated data for (a) α_1 (b) σ^2 (c) k_2 and (d) ϕ .




In the posterior analysis based on the Markov Chain Monte Carlo method, we need to satisfy two conditions in order to check if the posterior distribution is accurately described. The first is that the chain should become stationary. Figure 3.4 plots draws of the coefficient of the cumulative average rating, error variance, cutoff parameters and autocorrelation coefficients (for movie id 662) across iterations for Model 2. For this chain, movie fixed effect coefficients and autocorrelation parameters are started at 0, the coefficient of the average of the prior ratings is started at 1, the variance parameter σ^2 is started at 0.5 and the cutoff parameters are started at $(k_1, k_2, k_3) = (-1, 0.5, 1.5)$. The plots indicate that the Gibbs sampler reaches a stationary distribution rapidly.

If the starting values for cutoff parameters are too different from the true values, the algorithm takes significantly longer to converge. This is due to the conditional distribution for cutoff parameters, which allows smaller steps to be taken in each iteration as the sample

size increases. As suggested by Albert and Chib (1993), we proxy for the starting values of the cutoff parameters with the ones estimated by MLE with iid errors. The mean of the posterior draws in the stationary region provides the point estimates, and the standard deviation provides the standard errors.

Figure 3.4: Convergence plots using simulated data for (a) α_1 (b) σ^2 (c) k_2 and (d) ϕ for movie id 662.



Stationary Markov chains are necessary but not sufficient to gauge the performance of Gibbs sampler, because they may converge at a local mode and ignore other areas of higher posterior probability. One way to increase our confidence in the Gibbs sampler is to start chains from different regions in the parameter space. This will satisfy the second requirement if the estimated parameters are similar across different starting values.

We investigate the convergence of the Gibbs sampler by starting with different points in the parameter space. We compare the estimated posterior means and standard deviations for three chains. Chain 1 initializes the movie specific fixed effect coefficients and autocorrelation parameters at 0, coefficient of average of the prior ratings at 1 and the variance parameter at 0.5. Chain 2 initializes the movie specific fixed effect coefficients and the coefficient of the average of prior ratings at 0.5, autocorrelation parameters at 0.4 and the variance parameter at 1, while chain 3 initializes movie specific fixed effect coefficients

at 1, the coefficient of the average of the prior ratings at 0, autocorrelation parameters at 0.7 and the variance parameter at 0.1. The movie specific fixed effects are measured relative to the movie 'Inventing the Abbotts (1997)'. Movie-specific fixed effects capture the unobserved heterogeneity across movies.

The Gibbs sampler ran 75,000 iterations for each of the three chains while posterior means and standard deviations were calculated using the last 25,000 iterations. The estimated posterior means and standard deviations for the three chains are close to each other, providing evidence that the Markov chain generates draws from a region of the sample space with highest posterior probability. Results for chain 1 are presented in Table 3.1 and will be used for the analysis later. Although all the tests were done using the sample of 60 movies as described earlier, only the results for five movies are reported in Table 3.1 to save space.

Table 3.1: Posterior means and standard errors for Model 2 using three different Gibbs chains.

Parameters		Estimates (std. errors)
Intercepts	$\alpha_m(\text{movie id 662})$	0.41 (0.26)
	$\alpha_m(\text{movie id 1,642})$	2.15 (0.31)
	$\alpha_m(\text{movie id 1,788})$	1.59 (0.25)
	$\alpha_m(\text{movie id 2,136})$	0.79 (0.17)
	$\alpha_m(\text{movie id 2,380})$	1.19 (0.17)
Variable	α_1	1.18 (0.23)
Error Variance	σ^2	0.57 (0.10)
Autocorrelation coef.	$\phi_m(\text{movie id 662})$	0.23 (0.26)
	$\phi_m(\text{movie id 1,642})$	0.15 (0.30)
	$\phi_m(\text{movie id 1,788})$	-0.01 (0.25)
	$\phi_m(\text{movie id 2,136})$	0.21 (0.26)
	$\phi_m(\text{movie id 2,380})$	0.19 (0.36)
Cutoff Parameters	k_1	-1.45 (0.16)
	k_2	0.79 (0.09)
	k_3	2.08 (0.22)

Table 3.2 presents the posterior means and standard deviations of the fixed effects for Model 1 and Model 2. We note that the results of the two models are very close to one another.

This is possibly because Model 2 reveals that for only 5 movies in the sample have an autocorrelation coefficient significantly different from zero. This implies that for most movies there is no movie specific error that is autocorrelated over time. Figure 3.4(d) displays a typical convergence plot showing that autocorrelation parameter is not significantly different from zero.

We now turn our attention to the interpretation of α_1 . We get a positive and significant value of 1.18 using Model 2. The corresponding value of δ is 0.54. This means that individuals put roughly half of the weight on the average of the prior ratings and the other half on their private evaluations in order to arrive at their final ratings.

Table 3.2: Posterior mean and standard error of fixed effects for two models.


Parameters		Model 1	Model 2
Intercepts	$\alpha_m(\text{movie id 662})$	0.21 (0.15)	0.41 (0.26)
	$\alpha_m(\text{movie id 1,642})$	1.70 (0.24)	2.15 (0.31)
	$\alpha_m(\text{movie id 1,788})$	1.13 (0.22)	1.59 (0.25)
	$\alpha_m(\text{movie id 2,136})$	0.58 (0.15)	0.79 (0.17)
	$\alpha_m(\text{movie id 2,380})$	0.99 (0.15)	1.19 (0.17)
Variable	α_1	1.48 (0.18)	1.18 (0.23)

We also tested if depends on the time since release of the title. We found that α_1 decreases (p-value=0.08) as movies get older. This indicates that individuals may put lower weight on the average of the prior ratings for older movies. This is probably because they consider themselves relatively more informed about older movies than about recent releases.

3.5 Sensitivity analysis


We start with a question: Does the pattern of ratings change because of alpha not being equal to zero? We simulated ratings for 1,000 rating occasions using all the parameter estimates from Model 2. In an additional condition, we used all the estimates from Model 2 except the coefficient for the average of the prior ratings (α_1) which was set to zero. The pattern of the ratings obtained in the two conditions is shown in Figure 3.5. While the means of the ratings are very similar in the two conditions, the variance of the ratings in the first condition is smaller (0.90 compared to 2.22). The tendency of individuals to incorporate the information presented in the average ratings changed the pattern of ratings.

Figure 3.5: Simulated ratings showing the effect of incorporating cumulative average rating in evaluations.



We also checked the effect of receiving early positive or negative shocks to the ratings on the overall pattern of ratings. For example, this shock could be a result of ratings given in the beginning, with the intention of affecting subsequent ratings upwards or downwards. In the first case, we restrict five of the first 6 ratings to be equal to one. In the second case, we proceed in a similar fashion by restricting five of the first six ratings to be equal to five. We look at the resulting pattern of the first 1,000 ratings. Figure 3.6 shows the effect of early low or high ratings on the subsequent evaluations. The average of 1,000 ratings is 2.98 in the second case, compared to 2.72 in the first case. We note that the difference in the cumulative average ratings persists even after 1,000 ratings.

Figure 3.6: Simulated ratings showing the effect of early shock on subsequent evaluations.



3.6 Conclusion

The Internet has created new ways of gathering as well as communicating information. Anonymous online ratings are a quick and popular way to communicate information about a product by summarizing one's overall experience on an ordinal scale. This paper explores the process of reporting online ratings. In particular we examine if individuals incorporate the average of the prior ratings reported on the screen in their evaluations when they visit a particular website, in order to report their own experiences through anonymous online ratings.

We find the coefficient of cumulative average ratings to be significant and positive, suggesting that users incorporate the average of previous ratings displayed on screen in their own evaluations of the movie quality. The magnitude of the coefficient suggests that individuals put roughly half of the weight on the average of the prior ratings and the other half of the weight on their own evaluations in order to arrive at their final ratings. We can also infer that netflix.com users do not report ratings with the objective of influencing others' opinions, given the distribution of ratings we observe. We also find that for most movies the movie specific unobserved factors do not seem to be serially correlated. The value of the autocorrelation parameters was not significantly different from zero for 55 out of 60 movies.

Simulations using the estimated parameter values suggest that the tendency of individuals to incorporate the average rating in their own evaluations changes the pattern of the ratings significantly. We also find that early positive or negative shocks to the ratings get carried over to future rating occasions compared to the case when individuals ignore the average

ratings. This suggests that the introduction of a few high ratings at the start could significantly affect the ratings upwards. Hence, ratings appear to be susceptible to manipulation. Future work can focus on the inclusion of individual level and movie specific characteristics. In addition, trying to identify the mechanism responsible for the inclusion of the average ratings as information may be a worthwhile challenge for future research.

Bibliography

- V. Aguirregabiria and P. Mira. Sequential estimation of dynamic discrete games. *Econometrica*, 75(1):1–53, 2007. ISSN 00129682. URL <http://www.jstor.org/stable/4123107>.
- J. Albert and S. Chib. Bayesian analysis of binary and polychotomous response data. *Journal of American Statistical Association*, 88:669–679, 1993.
- G. Allenby and P. Lenk. Modeling household purchase behavior with logistic normal regression. *Journal of American Statistical Association*, 89:1218–1231, 1994.
- B. N. Anand and R. Shachar. (noisy) communication. *Quantitative Marketing and Economics*, 5(3):211–237, 2007.
- D. Ariely, G. Loewenstein, and D. Prelec. "coherent arbitrariness": Stable demand curves without stable preferences. *The Quarterly Journal of Economics*, 118(1):pp. 73–105, 2003. ISSN 00335533. URL <http://www.jstor.org/stable/25053899>.
- O. Armantier and O. Richard. Exchanges of cost information in the airline industry. *The RAND Journal of Economics*, 34(3):461–477, 2003. ISSN 07416261. URL <http://www.jstor.org/stable/1593741>.
- S. Athey and P. A. Haile. Empirical models of auctions. *Working Paper*, March 2006. URL <http://www.nber.org/papers/w12126>.
- K. Bagwell and G. Ramey. Advertising as information: Matching products to buyers. *Journal of Economics & Management Strategy*, 2(2):199–243, 1993.
- P. Bajari, C. L. Benkard, and J. Levin. Estimating dynamic models of imperfect competition. *Econometrica*, 75(5):1331–1370, 2007. ISSN 00129682. URL <http://www.jstor.org/stable/4502033>.
- P. Bajari, V. Chernozhukov, D. Nekipelov, and H. Hong. Semiparametric estimation of dynamic games of incomplete information. *Working Paper*, 2009.

- A. Banerjee. A simple model of herd behavior. *The Quarterly Journal of Economics*, 107: 797–817, 1992.
- G. Becker. A note on restaurant pricing and other examples of social influences on price. *The Journal of Political Economy*, 99:1109–1116, 1991.
- S. Bikhchandani, D. Hirshleifer, and I. Welch. Learning from behavior of others: Conformity, fads and information cascades. *The Journal of Economic Perspectives*, 12:151–170, 1998.
- G. E. P. Box and D. R. Cox. An analysis of transformations. *Journal of the Royal Statistical Society. Series B (Methodological)*, 26(2):211–252, 1964. ISSN 00359246. URL <http://www.jstor.org/stable/2984418>.
- T. Brock and H. Fromkin. Cognitive tuning set and behavioral receptivity to discrepant information. *Journal of Personality*, 36:108–125, 1968.
- R. E. Caves and D. P. Greene. Brands’ quality levels, prices, and advertising outlays: empirical evidence on signals and information costs. *International Journal of Industrial Organization*, 14(1):29–52, 1996. URL <http://ideas.repec.org/a/eee/indorg/v14y1996i1p29-52.html>.
- A. Chakraborty and R. Harbaugh. Clearly biased experts. *working paper*, 2008.
- Y. Chen, C. Narasimhan, and Z. J. Zhang. Individual marketing with imperfect targetability. *Marketing Science*, 20(1):23–41, 2001. ISSN 07322399. URL <http://www.jstor.org/stable/193220>.
- W. Chu. Demand signalling and screening in channels of distribution. *Marketing Science*, 11(4):327–347, 1992. ISSN 07322399. URL <http://www.jstor.org/stable/184092>.
- J. Chung and V. Rao. A general choice model for bundles with multiple-category products: Application to market segmentation and optimal pricing for bundles. *Journal of Marketing Research*, 40:115, 2003.
- R. N. Clarke. Collusion and the incentives for information sharing. *The Bell Journal of Economics*, 14(2):383–394, 1983. ISSN 0361915X. URL <http://www.jstor.org/stable/3003640>.
- A. Cohen. Cognitive tuning as a factor affecting impression formation. *Journal of Personality*, 29:235–245, 1961.
- V. P. Crawford and J. Sobel. Strategic information transmission. *Econometrica*, 50(6): 1431–1451, 1982. ISSN 00129682. URL <http://www.jstor.org/stable/1913390>.

- A. Cukierman. Asymmetric information and the electoral momentum of public opinion polls. *Public Choice*, 70:181–213, 1991.
- N. Epley and T. Gilovich. The anchoring-and-adjustment heuristic: Why the adjustments are insufficient. *Psychological science*, 17:311, 2006.
- C. Fershtman and A. Pakes. Finite state dynamic games with asymmetric information: A framework for applied work. *Working Paper*, 2009.
- E. Gal-Or. Information sharing in oligopoly. *Econometrica*, 53(2):329–343, 1985. ISSN 00129682. URL <http://www.jstor.org/stable/1911239>.
- E. Gal-Or. Information transmission—cournot and bertrand equilibria. *The Review of Economic Studies*, 53(1):85–92, 1986. ISSN 00346527. URL <http://www.jstor.org/stable/2297593>.
- A. Gelfand and A. Smith. Sampling-based approaches to calculating marginal densities. *Journal of American Statistical Association*, 85:398–409, 1990.
- ICE-Corporation. *Cost effective IC manufacturing 1998 to 1999*. Scottsdale, AZ: Integrated Circuit Engineering Corporation, 1997.
- K. Jacowitz and D. Kahneman. Measures of anchoring in estimation tasks. *Personality and Social Psychology Bulletin*, 21:1161–1166, 1995.
- J. S. Jedidi, K. and P. Manchanda. Measuring heterogeneous reservation prices for product bundles. *Marketing Science*, 22:107–130, 2003.
- J. Y. Jin. A comment on 'a general model of information sharing in oligopoly,' vol. 71 (1996), 260–288. *Journal of Economic Theory*, 93(1):144 – 145, 2000. ISSN 0022-0531. doi: DOI:10.1006/jeth.1999.2638. URL <http://www.sciencedirect.com/science/article/B6WJ3-45R86M1-7/2/99185a1ae56aa3f9e6a91e5a6cbe82b3>.
- B. Jullien. Participation constraints in adverse selection models. *Journal of Economic Theory*, 93(1):1 – 47, 2000. ISSN 0022-0531. doi: DOI:10.1006/jeth.1999.2641. URL <http://www.sciencedirect.com/science/article/B6WJ3-45R86M1-1/2/69ca8ddd58c8f7374ace81685a8d8896>.
- A. Kirmani and A. R. Rao. No pain, no gain: A critical review of the literature on signaling unobservable product quality. *The Journal of Marketing*, 64(2):66–79, 2000. ISSN 00222429. URL <http://www.jstor.org/stable/3203443>.
- J.-J. Laffont and D. Martimort. *The theory of incentives: the principal-agent model*. Princeton University Press, Princeton, N.J., 2002.

- M. A. Lariviere and V. Padmanabhan. Slotting allowances and new product introductions. *Marketing Science*, 16(2):112–128, 1997. ISSN 07322399. URL <http://www.jstor.org/stable/184214>.
- T. R. Lewis and D. E. M. Sappington. Countervailing incentives in agency problems. *Journal of Economic Theory*, 49(2):294 – 313, 1989. ISSN 0022-0531. doi: DOI: 10.1016/0022-0531(89)90083-5. URL <http://www.sciencedirect.com/science/article/B6WJ3-4CYG92V-43/2/bb41ba072c25fd47afc1a4f4b77a6e55>.
- G. Maggi and A. Rodriguez-Clare. On countervailing incentives. *Journal of Economic Theory*, 66(1):238 – 263, 1995. ISSN 0022-0531. doi: DOI:10.1006/jeth.1995.1040. URL <http://www.sciencedirect.com/science/article/B6WJ3-45S93BX-4D/2/9b009429d66ae03f5d07b9930d7af8c0>.
- A. A. Manchanda, P. and S. Gupta. The shopping basket: A model for multicategory purchase incidence decisions. *Marketing Science*, 18:95–114, 1999.
- R. McCulloch and P. Rossi. An exact likelihood analysis of the multinomial probit model. *Journal of Econometrics*, 64:207–240, 1994.
- P. Milgrom and J. Roberts. Price and advertising signals of product quality. *The Journal of Political Economy*, 94(4):796–821, 1986. ISSN 00223808. URL <http://www.jstor.org/stable/1833203>.
- P. R. Milgrom. Good news and bad news: Representation theorems and applications. *The Bell Journal of Economics*, 12(2):380–391, 1981. ISSN 0361915X. URL <http://www.jstor.org/stable/3003562>.
- L. E. Murillo. *The International Dynamic Random Access Memory Industry from 1970 to 1993 examined under the Dynamic Capabilities Prism: implications for technology policy*. PhD thesis, University of California, Berkeley, 1993.
- T. Mussweiler and F. Strack. The euro in the common european market: A single currency increases the comparability of prices. *Journal of Economic Psychology*, 25:557–563, 2004.
- P. Nelson. Information and consumer behavior. *The Journal of Political Economy*, 78(2): 311–329, 1970. ISSN 00223808. URL <http://www.jstor.org/stable/1830691>.
- W. Novshek and H. Sonnenschein. Fulfilled expectations cournot duopoly with information acquisition and release. *The Bell Journal of Economics*, 13(1):214–218, 1982. ISSN 0361915X. URL <http://www.jstor.org/stable/3003442>.

- M. Pesendorfer and P. Schmidt-Dengler. Identification and estimation of dynamic games. *Mimeo, London School of Economics*, 2003.
- S. Plous. Thinking the unthinkable: The effects of anchoring on likelihood estimates of nuclear war. *Journal of Applied Social Psychology*, 19:67–91, 1989.
- R. Radner. Team decision problems. *The Annals of Mathematical Statistics*, 33(3):857–881, 1962. ISSN 00034851. URL <http://www.jstor.org/stable/2237863>.
- M. Raith. A general model of information sharing in oligopoly. *Journal of Economic Theory*, 71(1):260 – 288, 1996. ISSN 0022-0531. doi: DOI:10.1006/jeth.1996.0117. URL <http://www.sciencedirect.com/science/article/B6WJ3-45NJKYF-29/2/17b702237b179d56c70fd3361bc970b6>.
- S. Ryan. The costs of environmental regulation in a concentrated industry. *Working Paper*, 2009.
- T. C. Schelling. *The strategy of conflict*. London: Oxford University Press, 1960.
- A. Schlosser. Posting versus lurking: Communication in a multiple audience context. *Journal of Consumer Research*, 32:260–265, 2005.
- C. Shapiro. Exchange of cost information in oligopoly. *The Review of Economic Studies*, 53(3):433–446, 1986. ISSN 00346527. URL <http://www.jstor.org/stable/2297638>.
- L. A. Stole. *Price Discrimination and Competition.*, volume 3 of *Handbook of Industrial Organization*, chapter 34, pages 2221–2299. Elsevier, 06 2007. URL <http://ideas.repec.org/h/eee/indchp/3-34.html>.
- F. Strack and T. Mussweiler. Explaining the enigmatic anchoring effect: Mechanism of selective accessibility. *Journal of Personality and Social Psychology*, 73:437–446, 1997.
- J. Tirole. *The Theory of Industrial Organization*. MIT Press, 1988.
- A. Tversky and D. Kahneman. Judgment under uncertainty: Heuristics and biases. *Science*, 185:1124–1131, 1974.
- J. M. Villas-Boas. Sleeping with the enemy: Should competitors share the same advertising agency? *Marketing Science*, 13(2):190–202, 1994. ISSN 07322399. URL <http://www.jstor.org/stable/183703>.
- J. M. Villas-Boas and U. Schmidt-Mohr. Oligopoly with asymmetric information: Differentiation in credit markets. *The RAND Journal of Economics*, 30(3):375–396, 1999. ISSN 07416261. URL <http://www.jstor.org/stable/2556054>.

- M. Villas-Boas. Product variety and endogenous pricing with evaluation costs. *Management Science* (forthcoming), 2009.
- X. Vives. Duopoly information equilibrium: Cournot and bertrand. *Journal of Economic Theory*, 34(1):71 – 94, 1984. ISSN 0022-0531. doi: DOI:10.1016/0022-0531(84)90162-5. URL <http://www.sciencedirect.com/science/article/B6WJ3-4CYGBT9-N6/2/ec2297bd600b510c14b458a41c2fe06d>.
- T. Yunogami. Mistake of japanese semiconductor industry. *AZojomo, Journal of Materials Online*, 2, 2006.
- J. Zhang. The sound of silence: Observational learning in the u.s. kidney market. *Marketing Science*, 29:315–335, 2010.

Appendix A

Appendix to Chapter 1

A.1 Posterior Distribution Given Demand Signal

The normality assumption provides a closed-form conditional distribution of the demand shock and other firms' signals. In particular the conditional distribution of the demand shock and of the competitors of firm 1 (without loss of generality) and of the demand shock given its own private demand signal, s_1 , are distributed according to

$$\begin{matrix} \varepsilon \\ s_2 \\ \vdots \\ s_N \end{matrix} \bigg| s_1 \sim N \left(\begin{bmatrix} \frac{\sigma_\varepsilon^2}{\sigma_\varepsilon^2 + \sigma_\xi^2 + \sigma_\eta^2} s_1 \\ \frac{\sigma_\varepsilon^2 + \sigma_\eta^2}{\sigma_\varepsilon^2 + \sigma_\xi^2 + \sigma_\eta^2} s_1 \\ \vdots \\ \frac{\sigma_\varepsilon^2 + \sigma_\eta^2}{\sigma_\varepsilon^2 + \sigma_\xi^2 + \sigma_\eta^2} s_1 \end{bmatrix}, \begin{bmatrix} \left(\frac{1}{\sigma_\varepsilon^2} + \frac{1}{\sigma_\xi^2 + \sigma_\eta^2} \right)^{-1} & \frac{\sigma_\varepsilon^2 \sigma_\xi^2}{\sigma_\varepsilon^2 + \sigma_\xi^2 + \sigma_\eta^2} & \cdots & \frac{\sigma_\varepsilon^2 \sigma_\xi^2}{\sigma_\varepsilon^2 + \sigma_\xi^2 + \sigma_\eta^2} \\ \frac{\sigma_\varepsilon^2 \sigma_\xi^2}{\sigma_\varepsilon^2 + \sigma_\xi^2 + \sigma_\eta^2} & 2\sigma_\xi^2 - \frac{\sigma_\xi^4}{\sigma_\varepsilon^2 + \sigma_\xi^2 + \sigma_\eta^2} & \cdots & \sigma_\xi^2 - \frac{\sigma_\xi^4}{\sigma_\varepsilon^2 + \sigma_\xi^2 + \sigma_\eta^2} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\sigma_\varepsilon^2 \sigma_\xi^2}{\sigma_\varepsilon^2 + \sigma_\xi^2 + \sigma_\eta^2} & \sigma_\xi^2 - \frac{\sigma_\xi^4}{\sigma_\varepsilon^2 + \sigma_\xi^2 + \sigma_\eta^2} & \cdots & 2\sigma_\xi^2 - \frac{\sigma_\xi^4}{\sigma_\varepsilon^2 + \sigma_\xi^2 + \sigma_\eta^2} \end{bmatrix} \right) \quad (\text{A.1})$$

A.2 Polynomial approximation of the Value Function

The *left hand side* of (1.16) is approximated by the polynomial expression:

$$\begin{aligned} V_i(\Omega_{t+1}, \Theta_{it+2}) \simeq P_i(\Omega_{t+2}, \Theta_{it+2}) &= \alpha_0 + \alpha_1 c_{it+1} + \alpha_2 \varepsilon_{t+1} + \alpha_3 K_{it+1} + \alpha_4 s_{it+2} \\ &+ \alpha_5 \sum_{-i} K_{-it+1} + \alpha_6 \sum_{-i} c_{-it+1} + \alpha_7 K_{it+1}^2 + \alpha_8 K_{it+1} s_{it+2} \\ &+ \alpha_9 \gamma^+ + \alpha_{10} \gamma^- + \alpha_{11} K_{it+1} c_{it+1} + \alpha_{12} K_{it+1} \sum_{-i} K_{-it+1} \\ &+ \alpha_{13} K_{it+1} \varepsilon_{t+1} \end{aligned}$$

A.3 First-Order Condition for Capacity Decisions

The first-order condition approximation for capacity decisions resulting from the polynomial $P_i(\Omega_{t+2}, \Theta_{it+2})$ above is given by:

$$\begin{aligned}
 FOC_{it} : \quad & \alpha_3 + 2\alpha_7 K_{it+2} + \alpha_8 \underbrace{E[s_{it+3}]}_{=0} + \alpha_{11} c_{it+2} \\
 & + \alpha_{12} \sum_{-i} E[K_{-it+2}(s_{-it+2})|s_{it+2}] + \\
 & + \alpha_{13} E[\epsilon_{t+2}|s_{it+2}] \\
 & + 1(K_{it+2} > K_{it+1})2\gamma^+(K_{it+2} - K_{it+1}) + 1(K_{it+2} < K_{it+1})2\gamma^-(K_{it+1} - K_{it+2}) = 0
 \end{aligned}$$

Note that this specification implies multiple equilibria in capacities. However, the equilibria are recoverable in the historic data, and their nature is assumed to be unchanged in policy simulations. In addition its solution is defined only if $\alpha_7 \neq \gamma^+$ and $\alpha_7 \neq \gamma^-$.

A.4 Dynamic Parameter Estimates

Table A.1: Dynamic Parameter Estimates

Parameter	Coefficient	Std. Error	Parameter	Coefficient	Std. Error
α_0	-0.736	—	α_7	-0.515	—
α_1	6.095	—	α_8	4.633	—
α_2	-1.809	—	α_9	5.551	—
α_3	-3.655	—	α_{10}	-4.914	—
α_4	9.720	—	α_{11}	-2.666	—
α_5	-1.968	—	α_{12}	0.352	—
α_6	-3.456	—	α_{13}	4.164	—

Appendix B

Appendix to Chapter 2

As discussed in Section 2.2, a product quality range is imposed such that the firm is better off serving both types of consumers rather than just serving one, and also such that all price equilibria are interior. Let such a range be $[\underline{q}, \bar{q}]$, such that $q_L > \underline{q}$ and $q_H < \bar{q}$. I equate the profits of serving the low (high) customer segment to those of serving both customers in order to find $\underline{q}'$ ($\bar{q}'$) while assuming that the optimal prices are interior at such quality levels. I then check for the conditions for interior price solutions and identify the non-redundant ones. Throughout this analysis it is assumed that $\gamma > t_H$, such that if any consumer faces a product of low (e.g.: zero) quality, she is still willing to take it at zero cost as long as she gets the highest possible utility draw. This condition is not required but when assumed greatly simplifies the relevant parameter space for q and it does not affect any of the results below. The general conditions are available from the author.

B.1 Quality Range

The auxiliary lower bound for product quality $\underline{q}'$ solves the expression:

$$\begin{aligned} \pi_{\underline{q}'}(L) &= \pi_{\underline{q}'}(L, H) \\ \Leftrightarrow \frac{(\theta_L \underline{q}' + \gamma - t_L)^2}{8\gamma} &= \frac{((\theta_H + \theta_L) \underline{q}' - t_H - t_L + 2\gamma)^2}{16\gamma} \end{aligned} \quad (\text{B.1})$$

$$\Leftrightarrow \underline{q}' = \frac{t_H - (\sqrt{2} - 1)t_L - (2 - \sqrt{2})\gamma}{\theta_H - (\sqrt{2} - 1)\theta_L} \quad (\text{B.2})$$

The auxiliary upper bound quality level $\bar{q}'$ solves:

$$\begin{aligned} \pi_{\bar{q}'}(H) &= \pi_{\bar{q}'}(L, H) \\ \Leftrightarrow \frac{(\theta_H \bar{q}' + \gamma - t_H)^2}{8\gamma} &= \frac{((\theta_H + \theta_L) \bar{q}' - t_H - t_L + 2\gamma)^2}{16\gamma} \end{aligned} \quad (\text{B.3})$$

$$\Leftrightarrow \bar{q}' = \frac{(\sqrt{2} - 1)t_H - t_L + (2 - \sqrt{2})\gamma}{(\sqrt{2} - 1)\theta_H - \theta_L} \quad (\text{B.4})$$

where the profit expressions come directly from substituting the optimal price from equation (2.5) into (2.4), and the analogue procedure is done for the case where the firm serves both consumer segments. When

$$\theta_H > \frac{\theta_L}{(\sqrt{2} - 1)},$$

$\bar{q}'$ is positive such that over this quality threshold the firm prefers to price the consumers of low type out of the market. Notice that this implies $\theta_H > 2\theta_L$.

B.2 Interior Profit Maximization

Single-customer segment:

By taking the first order condition of equation (2.4), in the scenario of interior solutions, and solving it with respect to price, we get that p^* is interior for each of the consumer segments if and only if:

$$0 \leq \frac{\theta_S q_j - t_S + \gamma}{4\gamma} \leq 1$$

for $\forall S \in \{L, H\}, \forall j \in \{L, M, H\}$. Given the non-negativity conditions on the remaining parameters and the assumption $\gamma > t_H$, the set of equations above can be reduced to the equation:

$$q < \frac{3\gamma + t_H}{\theta_H} \quad (\text{B.5})$$

Two customer segments:

Let $p_L^*(p_H^*)$ maximize the firm's profits when only the low (high) type consumers visit. Then in order to ensure that all outcomes are interior for when consumers deviate (i.e.: both consumers visit), one needs:

$$0 \leq \frac{(2\theta_L - \theta_H)q_j + \gamma + t_H - 2t_L}{4\gamma} \leq 1$$

$$0 \leq \frac{(2\theta_H - \theta_L)q_j + \gamma + t_L - 2t_H}{4\gamma} \leq 1$$

$\forall I \in \{L, H\}, \forall j \in \{L, M, H\}$. These conditions can be summarized by:

$$q < \text{Min} \left\{ \frac{2t_H - t_L + 3\gamma}{2\theta_H - \theta_L}, \frac{t_H - 2t_L + \gamma}{\theta_H - 2\theta_L} \right\} \quad (\text{B.6})$$

and

$$q > \text{Max} \left\{ \frac{2t_H - t_L - \gamma}{2\theta_H - \theta_L}, \frac{t_H - 2t_L - 3\gamma}{\theta_H - 2\theta_L} \right\} \quad (\text{B.7})$$

The intuition behind these conditions is that quality is imposed not to be too high nor too low in relation to the support of the uncertainty component and in relation to the outside options, so that there is an interior solution. Otherwise the firm would like to price one of the consumer segments out of the market (for example when t_H is too high when compared to q , the firm could choose to only serve low type consumers), or alternatively would like to serve all consumers of one type: for example, when t_L is sufficiently low when compared to q_L , the firm would like to serve all low type consumers. Notice that all conditions are satisfied by allowing for a big enough γ . This is because as the option value increases with uncertainty, the quality of the product becomes less important in the expected utility of searching.

Interior prices and the limits of the quality range

Let $A \equiv \{\gamma > t_H > t_L > 0 \wedge \theta_H > 2\theta_L > 0\}$. Under A it is possible to check that we only need:

$$\underline{q} \equiv \text{Max} \left\{ 0, \frac{2t_H - t_L - \gamma}{2\theta_H - \theta_L} \right\} < q < \text{Min} \left\{ \frac{2t_H - t_L + 3\gamma}{2\theta_H - \theta_L}, \frac{t_H - 2t_L + \gamma}{\theta_H - 2\theta_L} \right\} \equiv \bar{q} \quad (\text{B.8})$$

The auxiliary limits $\underline{q}', \bar{q}'$ from (B.2) and (B.4) as well as condition (B.5) are redundant given condition (B.8).

B.3 Proof of Proposition 1

Expressions for ex-ante consumer surplus (2.1) and optimal price (2.5) yield the following net ex-ante expected consumer surplus of searching conditional on quality level q :

$$CS_S(q|p_S^*) \equiv \frac{(\theta_S q - t_S + \gamma)^2}{16\gamma} - k$$

where p_S^* is defined as the profit-maximizing price for firm with quality q facing only consumers of type S . It follows that

$$\frac{\partial CS_S(q|p_S^*)}{\partial q} = \frac{\theta_S(\theta_S q - t_S + \gamma)}{8\gamma} > 0$$

by $\gamma > t_H$. In addition the comparative statics on θ_S, t_S become:

$$\frac{\partial CS_S(q|p_S^*)}{\partial \theta_S} = \frac{q(\theta_S q - t_S + \gamma)}{8\gamma} > 0$$

$$\frac{\partial CS_S(q|p_S^*)}{\partial t_S} = -\frac{(\theta_S q - t_S + \gamma)}{8\gamma} < 0$$

B.4 Proof of Proposition 2

The ex-ante consumer surplus of the high-type customers given that they face p_L^* (where p_L^* maximizes the firm's profits when only the low type consumers visit) is:

$$CS_H(q|p_L^*) = \frac{[(2\theta_H - \theta_L)q + \gamma - 2t_H + t_L]^2}{16\gamma}$$

and the analogous consumer surplus for the low type consumers is:

$$CS_L(q|p_H^*) = \frac{[(2\theta_L - \theta_H)q + \gamma - 2t_L + t_H]^2}{16\gamma}$$

It is easy to check that

$$\frac{\partial CS_H(p_L^*) - k}{\partial q} > 0$$

and

$$\frac{\partial CS_L(p_H^*) - k}{\partial q} < 0$$

whenever the interior condition (B.8) holds and $\theta_H > 2\theta_L$.

B.5 Proof of Proposition 3

Let $p_{S'}^*$ be the equilibrium price for the firm given visits by consumer segments in set $S' \equiv \{L, H, LH\}$. By comparison of the consumer surplus expression (2.1), the following result arises:

$$\begin{cases} CS_L(q|p_L^*) < CS_L(q|p_{LH}^*) \\ CS_H(q|p_H^*) > CS_H(q|p_{LH}^*) \end{cases} \quad \text{if } q < \tilde{q}$$

$$\begin{cases} CS_L(q|p_L^*) > CS_L(q|p_{LH}^*) \\ CS_H(q|p_H^*) < CS_H(q|p_{LH}^*) \end{cases} \quad \text{if } q > \tilde{q}$$

B.6 Definition of $q_{(1)}$ to $q_{(3)}$

Each of these quality thresholds is defined by the intersection of k and different expected consumer surpluses:

$$\begin{aligned} CS_L(q_{(1)}|p_L^*) &= k \\ CS_H(q_{(2)}|p_L^*) &= k \\ CS_L(q_{(3)}|p_H^*) &= k \end{aligned}$$

B.7 Equilibria with Matching Interests

Let $q_L, q_M \in [q_{(1)}, q_{(2)}]$, $q_H > q_{(3)}$, firms q_L, q_M set price p_L^* and firm q_H set price p_H^* . Informativeness holds by the stated assumptions which imply:

$$CS_L(\{q_L, q_M\}) > k \tag{B.9}$$

$$CS_H(\{q_H\}) > k \quad (\text{B.10})$$

For decisiveness, we need one of the two conditions to hold:

$$CS_L(\{q_L, q_M, q_H\}) < k \quad (\text{B.11})$$

$$CS_H(\{q_L, q_M, q_H\}) < k \quad (\text{B.12})$$

The existence of such equilibria is intuitive and given the closed form of the model, is easily shown by algebraic substitution. The intuition for each is also simple: since in the case of the low-type consumers the ‘decisive’ difference between expressions (B.9) and (B.11) comes through $CS_L(q_H)$, and additionally since we know that such surplus is decreasing in q_H whenever $\theta_H > 2\theta_L$, it suffices for q_H to be big enough (in relation to q_L and q_M) in order for the information to be decisive for the low-type consumers. The equilibrium may also be decisive for the high type consumers: in the case where q_H is small enough in relation to q_L and q_M by the same type of reasoning. Finally, the intuition behind the other type of matching equilibrium described in Section 2.3.2 follows a similar reasoning.

B.8 Incentive to Misreport

In this case, $q_L < q_{(1)}$, $q_M \in [q_{(1)}, q_{(2)}]$, $q_H > q_{(3)}$, where firms q_L, q_M set price p_L^* , and firm q_H sets price p_H^* . In that case informativeness holds by the assumptions, which imply:

$$CS_L(q_L) < k \quad (\text{B.13})$$

$$CS_L(q_M) > k \quad (\text{B.14})$$

$$CS_H(q_H) > k \quad (\text{B.15})$$

For decisiveness, we need one of the two conditions to hold:

$$CS_L(\{q_L, q_M, q_H\}) < k \quad (\text{B.16})$$

$$CS_H(\{q_L, q_M, q_H\}) < k \quad (\text{B.17})$$

and in addition we need that:

$$CS_L(\{q_L, q_M\}) > k \quad (\text{B.18})$$

such that the low type consumers search in equilibrium. This happens whenever q_M is high relative to q_L . Finally, condition (B.16) holds when q_L is low and q_H is high relative to q_M whereas condition (B.17) holds when q_H is sufficiently low relative to q_L and q_M .

B.9 Continuous Quality

This section describes how a model of continuous quality does not imply that there will exist a product quality level that leads all consumers to search, and hence communication is still possible in such a case. Assume quality follows a continuous distribution function F_q with support $[q, \bar{q}]$, and let the consumer marginal valuations follow a continuous distribution function F_θ . Finally, assume a continuous outside option function $t(\theta)$ such that $t'(\theta) > 0$ and $t''(\theta) > 0$. In that case, consumer i buys the product after search iff $\theta_i q + \varepsilon_i - p > t(\theta_i)$. In addition, she visits the store iff $E(u_i | \hat{q}) > k + t(\theta)$. Using an analogous formula to 2.1, we get that consumer i searches under degenerate beliefs $\hat{q}$ iff:

$$\theta_i \hat{q} - p^*(\hat{q}) + \gamma - t(\theta_i) > 2\sqrt{k\gamma} \quad (\text{B.19})$$

Given the properties of $t(\theta)$, and an interior equilibrium price $p^*(\hat{q})$, one can find the threshold consumers who are willing to visit. Let $\bar{\theta}$ and $\underline{\theta}$ be such thresholds, which solve the equilibrium condition B.19 under equality (notice that the convex nature of $t(\theta)$ provides two solutions to that equation). Then, under beliefs $\hat{q}$, only consumers with $\theta \in [\underline{\theta}, \bar{\theta}]$ are willing to search. Assume $[\underline{\theta}, \bar{\theta}]$ and $p^*(\hat{q})$ are well-defined and unique under all the support of F_q .

Finally, notice that the impact of increments of quality on the expected consumer surplus can be decomposed into the direct quality effect, and the indirect pricing effect:

$$\frac{dE(u_i)}{dq} = \frac{\partial E(u_i)}{\partial q} + \frac{\partial E(u_i)}{\partial p^*} \frac{dp^*}{dq}$$

When the first term dominates the second term, more quality is good news for consumers, while when the second term dominates the first, the indirect pricing effect dominates the quality effect. Sufficient conditions for the absence of a “perfect product” that could attract all consumers to search are:

$$\begin{aligned} \left. \frac{\partial E(u_i)}{\partial q} \right|_{\theta_i = \bar{\theta}} &> \left. \frac{\partial E(u_i)}{\partial p^*} \frac{dp^*}{dq} \right|_{\theta_i = \bar{\theta}} \\ \left. \frac{\partial E(u_i)}{\partial q} \right|_{\theta_i = \underline{\theta}} &< \left. \frac{\partial E(u_i)}{\partial p^*} \frac{dp^*}{dq} \right|_{\theta_i = \underline{\theta}} \end{aligned}$$

The first condition says that more quality is good news for consumers who are near the upper region of types currently willing to search under quality/belief q . The second condition says that more quality is bad news for consumers near the end of the distribution of types who are willing to visit under quality/belief q . In this sense, a slight increase in quality attracts more types above, but less types below. When these inequalities hold for all the support of F_q , and the function $E(u_i)$ is continuous on type θ_i , then more quality

always attracts some consumers and scares away others, which leads us to the intuition that no single quality level will attract all consumer types. In this case, communication carries the necessary trade-off among consumer segments that enables informative advertising to take place.

Appendix C

Appendix to Chapter 3

The Gibbs sampler proceeds by drawing iteratively from the conditional distributions described below, in the same order as they are presented.

C.1 Generating Y

The posterior distribution for Y given the data and other parameters is

$$[Y, R, \varepsilon_{\cdot,0}, \alpha, \Phi, \sigma^2, K] \propto [R|K, Y][Y|\varepsilon_{\cdot,0}, \alpha, \Phi, \sigma^2]$$

Because $\{y_{m,t}\}$ are conditionally independent for different movies, we can generate $y_{m,t}$ independently for each movie. Within a movie, $y_{\cdot,t}$ have a hierarchical structure due to their autocorrelation:

$$[y_{m,\cdot}|\varepsilon_{m,0}] = [y_{m,1}|\varepsilon_{m,0}] \prod_{t=2}^{T_m} [y_{m,t}|y_{m,t-1}].$$

The vector Y is sequentially generated from $y_{m,1}$ to y_{m,T_m} .

Define

$$\mu_{m,t} \equiv X_{m,t} \alpha; e_{m,0} \equiv \varepsilon_{m,0} \tag{C.1}$$

and

$$e_{m,t} = y_{m,t} - \mu_{m,t} \text{ for } t = 1, \dots, T_m \tag{C.2}$$

where α and $y_{m,t}$ are the most recent random draws from the Gibbs sampler.

The conditional distribution of $y_{m,t}$ for $t = 1, \dots, T_{m-1}$ is

$$\begin{aligned}
[y_{m,t} | r_{m,\cdot}, y_{m,1}, \dots, y_{m,t-1}, y_{m,t+1}, \dots, y_{m,T_m}] &\propto p_{m,t}(r_{m,t}) [y_{m,t} | y_{m,t-1}] [y_{m,t+1} | y_{m,t}] \\
&\propto p_{m,t}(r_{m,t}) \exp\left\{-\frac{1}{2}(e_{m,t} - v_{m,t})^2 V_m^{-1}\right\} \quad (C.3)
\end{aligned}$$

where

$$v_{m,t} = V_m \frac{\phi_m}{\sigma^2} (e_{m,t-1} + e_{m,t+1})$$

and

$$V_m = \frac{\sigma^2}{1 + \phi_m^2}.$$

For the last period, T_m , the conditional distribution is

$$\begin{aligned}
[y_{m,T_m} | r_{m,\cdot}, y_{m,1}, \dots, y_{m,T_m-1}] &\propto p_{m,T_m}(r_{m,T_m}) \exp\left\{-\frac{1}{2\sigma^2} (e_{m,T_m} - \phi_m e_{m,T_m-1})^2\right\} \\
&\propto p_{m,t}(r_{m,t}) \quad (C.4)
\end{aligned}$$

Rejection sampling is used to generate $y_{m,t}$. Starting a period $t = 1$, a candidate for $y_{m,t}$ is generated from the normal density in equation (C.3) or (C.4), and the test function $p_{m,t}(r_{m,t})$ is computed. If the test function is greater than a uniform random number, then the candidate is accepted and the residual, $e_{m,t}$, is updated. Then t is increased by one until $y_{m,1}$ to y_{m,T_m} are generated. When generating $y_{m,t}$ for next period the updated $e_{m,t}$ from previous period is used. Logistic errors ($\xi_{m,t}$) are also generated corresponding to each $y_{m,t}$, and saved for generating K later.

C.2 Generating $\varepsilon_{\cdot,0}$

The posterior distribution of $\varepsilon_{\cdot,0}$ given the other parameters is

$$[\varepsilon_{\cdot,0} | R, Y, \alpha, \Phi, \sigma^2, K] \propto \exp\left\{-\frac{1}{2c^2} \varepsilon_{m,0}^2\right\} \exp\left\{-\frac{\phi_m^2}{2\sigma^2} \left(\varepsilon_{m,0} - \frac{e_{m,1}}{\phi_m}\right)^2\right\}$$

where $e_{m,1}$ is defined by equations (C.1) and (C.2), and $c^2 = 1$. $\varepsilon_{m,0}$ is generated from $N(u, V)$ where

$$u = \frac{V \phi_m e_{m,1}}{\sigma^2}$$

and

$$V = \left(\frac{\phi_m^2}{\sigma^2} + \frac{1}{c^2} \right)^{-1}.$$

C.3 Generating α

The posterior distribution of α given the remaining parameters is

$$[\alpha | R, Y, \varepsilon_{\cdot,0}, \alpha, \Phi, \sigma^2, K] \propto [Y | \varepsilon_{\cdot,0}, \alpha, \Phi, \sigma^2][\alpha].$$

For movie m let

$$y_{m,1}^* = y_{m,1} - \phi_m \varepsilon_{m,0}$$

and

$$X_{m,1}^* = X_{m,1}$$

for the first period, and

$$y_{m,t}^* = y_{m,t} - \phi_m y_{m,t-1}$$

and

$$X_{m,t}^* = X_{m,t} - \phi_m X_{m,t-1}$$

for $t = 2, 3, \dots, T_m$.

The likelihood of α is proportional to

$$\exp \left\{ -\frac{1}{2\sigma^2} (Y^* - X^* \alpha)' (Y^* - X^* \alpha) \right\}$$

which gives rise to Zellner's seemingly unrelated regression equation (SURE) model. If the prior for α is $N_p(a_0, A_0)$, then we generate α from $N_p(a, A)$ where p is equal to the dimension of α and

$$A = \left(\frac{1}{\sigma^2} X^{*'} X^* + A_0^{-1} \right)^{-1}$$

and

$$a = A \left(\frac{1}{\sigma^2} X^{*'} X^* \hat{a} + A_0^{-1} a_0 \right)$$

where

$$\hat{a} = \left(\frac{1}{\sigma^2} X^{*'} X^* \right)^{-1} \left(\frac{1}{\sigma^2} X^{*'} Y \right).$$

If the prior for α is non-informative, then α can be generated from $N_p \left(\hat{a}, \left[\frac{1}{\sigma^2} X^{*'} X^* \right]^{-1} \right)$.

C.4 Generating Φ

The posterior distribution of Φ given the other parameters is

$$\begin{aligned} [\Phi | R, Y, \varepsilon_{\cdot,0}, \alpha, \sigma^2, K] &\propto [Y | \varepsilon_{\cdot,0}, \alpha, \Phi, \sigma^2] [\Phi] \\ &\propto \exp \left\{ -\frac{1}{2\sigma^2} \sum_{m=1}^M \sum_{t=1}^{T_m} (e_{m,t} - \phi_m e_{m,t-1})^2 \right\} [\Phi] \\ &\propto \exp \left\{ -\frac{1}{2} (\phi - \theta)' W (\phi - \theta) \right\} \prod_{j=1}^M I(-1 < \phi(j) < 1) \end{aligned}$$

where

$$W = \frac{1}{\sigma^2} \begin{bmatrix} \sum_t e_{1,t-1}^2 & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & \sum_t e_{M,t-1}^2 \end{bmatrix}$$

and

$$\theta = W^{-1} \begin{bmatrix} \sum_t e_{1,t} \cdot e_{1,t-1} \\ \vdots \\ \sum_t e_{M,t} \cdot e_{M,t-1} \end{bmatrix}.$$

Finally,

$$e_{m,0} = \varepsilon_{m,0}$$

and

$$e_{m,t} = y_{m,t} - X_{m,t} \alpha \text{ for } t = 1, \dots, T_m.$$

The prior for ϕ_m is given by $U(-1, 1)$.

As a result, Φ is generated from a truncated multivariate normal distribution with support $[-1, 1]^M$. (That is, we keep drawing from the distribution until all the coordinated are between -1 and 1.)

C.5 Generating σ^2

The posterior distribution of σ^2 given the other parameters is

$$\begin{aligned} [\sigma^2 | R, Y, \varepsilon_{\cdot,0}, \alpha, \Phi, \sigma^2, K] &\propto [Y | \varepsilon_{\cdot,0}, \alpha, \Phi, \sigma^2] [\sigma^2] \\ &\propto \sigma^{-N/2} \exp \left\{ -\frac{1}{\sigma^2} \sum_m \sum_t e_{m,t}^2 \right\} [\sigma^2] \end{aligned}$$

where

$$e_{m,0} = \varepsilon_{m,0}$$

$$e_{m,1} = y_{m,1} - X_{m,1} \alpha - \phi_m \varepsilon_{m,1}$$

$$e_{m,t} = (y_{m,t} - \phi_m y_{m,t-1}) - (X_{m,t} - \phi_m X_{m,t-1}) \text{ for } t = 2, \dots, T_m$$

and N is the total number of observations in the dataset.

The prior distribution of σ^2 is $IG(a_0, b_0)$ where $a_0 = 3$ and $b_0 = 0.001$ (as in Allenby and Rossi, 1999). Thus, σ^2 is generated from the posterior density $IG(a_1, b_1)$ where

$$a_1 = a_0 + \frac{N}{2}$$

$$b_1 = b_0 + \sum_{m=1}^M \sum_{t=1}^{T_m} e_{m,t}^2.$$

C.6 Generating K

The posterior distribution of k_j given the remaining parameters is

$$\begin{aligned} &[k_j | R, Y, \varepsilon_{\cdot,0}, \alpha, \Phi, \sigma^2] \\ &\propto \prod_{m=1}^M \prod_{t=1}^{T_m} I(r_{m,t} = j) (k_{j-1} < y_{m,t}^* < k_j) + I(r_{m,t} = j+1) (k_j < y_{m,t}^* < k_{j+1}). \end{aligned}$$

This conditional distribution is equivalent to the uniform distribution on the interval

$$[\max\{\max\{y_{m,t}^* : r_{m,t} = j\}, k_{j-1}\}, \min\{\min\{y_{m,t}^* : r_{m,t} = j+1\}, k_{j+1}\}]$$

where

$$y_{m,t}^* = y_{m,t} + \zeta_{m,t}.$$

Consequently, k_j is generated from $U(m, n)$ where $m = \max\{\max\{y_{m,t}^* : r_{m,t} = j\}, k_{j-1}\}$ and $n = \min\{\min\{y_{m,t}^* : r_{m,t} = j+1\}, k_{j+1}\}$.