

UC San Diego

UC San Diego Electronic Theses and Dissertations

Title

Essays on unconventional pricing strategies and impacts of economic regulation on stock return asymmetry

Permalink

<https://escholarship.org/uc/item/5j62531d>

Author

Lima, Daniel F.

Publication Date

2011

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA, SAN DIEGO

**Essays on Unconventional Pricing Strategies and Impacts of Economic
Regulation on Stock Return Asymmetry**

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

in

Economics

by

Daniel F. Lima

Committee in charge:

Professor Silke Forbes, Chair
Professor Karsten Hansen
Professor Ivana Komunjer
Professor Dominique Lauga
Professor David Miller

2011

Copyright
Daniel F. Lima, 2011
All rights reserved.

The dissertation of Daniel F. Lima is approved, and it is acceptable in quality and form for publication on microfilm and electronically:

Chair

University of California, San Diego

2011

DEDICATION

I dedicate this dissertation to my parents, Ailton and Marivone,
and my sister, Jane.

EPIGRAPH

*“Dignity consists not in possessing honors,
but in the consciousness that we deserve them.”*

— Aristotle

TABLE OF CONTENTS

Signature Page		iii
Dedication		iv
Epigraph		v
Table of Contents		vi
List of Figures		viii
List of Tables		ix
Acknowledgements		x
Vita		xi
Abstract of the Dissertation		xii
Chapter 1	Customer Feedback and Planned Obsolescence of High-Tech Products	1
	1.1 Introduction	2
	1.2 Information Acquisition <i>vs.</i> Reputation Damage	5
	1.3 The Model	8
	1.3.1 The Demand	9
	1.3.2 The Supply	11
	1.3.3 Equilibrium Analysis and Comparative Statics	14
	1.4 Case Study: Apples’s iPhone Release Strategy	18
	1.5 Conclusion	20
	1.6 Tables and Figures	22
Chapter 2	Product Downsizing and Hidden Price Changes in the Ready-to-Eat Cereal Market	25
	2.1 Introduction	26
	2.2 Hidden Price Changes	28
	2.2.1 Cost as a Driver of Downsizing	29
	2.2.2 Previous Research on Downsizing	31
	2.3 The Model	36
	2.3.1 Extending the AIDS Model	37
	2.3.2 The Supply	41
	2.3.3 Partial Equilibrium Analysis	42
	2.4 Data	45
	2.4.1 Identifying Downsizing	47

2.5	Empirical Results	49
2.5.1	Linear Demand	49
2.5.2	AIDS Predictions	52
2.5.3	A Diff-in-Diff Approach	54
2.5.4	Logit Demand and Profitability	57
2.6	Conclusions	58
2.7	Tables and Figures	60
Chapter 3	Economic Regulation and Asymmetry in Ex-Post Stock Returns	78
3.1	Introduction	79
3.2	Measuring Asymmetry in Financial Returns	83
3.2.1	Unconditional Asymmetry Tests	89
3.2.2	Conditional Asymmetry Tests	91
3.2.3	Empirical Results of Asymmetry Tests	92
3.3	The Relation Between Asymmetry and Regulatory Status	94
3.3.1	Contingency Tables	94
3.3.2	Equality of the Conditional Densities	97
3.3.3	Location and Scale Differences	99
3.4	Conclusions	101
3.5	Tables and Figures	102
Appendix A	Proof for Proposition 2.3.1	114
Appendix B	Moments of the Return Process	117
Appendix C	The Skewed-t Distribution	120
Bibliography	122

LIST OF FIGURES

Figure 1.1: iPhone Sales Chart	22
Figure 2.1: Input Commodity Prices (January 2006 = 100)	61
Figure 2.2: Aggregate Market Share	62
Figure 2.3: Aggregate Price Per Pound	63
Figure 2.4: Aggregate Price Per Box	64
Figure 2.5: General Mills Cheerios 15oz. to 14oz.	65
Figure 2.6: Possible Substitution Patterns Resulting from Downsizing	66
Figure C.1: Skewed versions of a Student-t Density	121

LIST OF TABLES

Table 1.1: Consumer Surplus	23
Table 1.2: Conditions Supporting Proposition 1 (grouped by players)	24
Table 2.1: Commodity Prices	61
Table 2.2: Aggregate Market Shares and Average Prices	67
Table 2.3: General Mills Downsizing	68
Table 2.4: General Mills Downsizing, Table 2.3 Continued	69
Table 2.5: Single Price Linear Demand - GK's Specification	70
Table 2.6: Full Specification, Linear Demand	71
Table 2.7: AIDS Prediction of Expenditure Share	72
Table 2.8: AIDS Prediction of Expenditure Share, Sorted by Ratio	73
Table 2.9: AIDS Predictions, Multiple Sizes of Same Brand	74
Table 2.10: Diff-in-Diff Expenditure Share	75
Table 2.11: Diff-in-Diff Expenditure Share, Table 2.10 Continued	76
Table 2.12: Price-Cost Margins	77
Table 3.1: Firms with significant skewness by the (3.9) and (3.11) tests. . .	103
Table 3.2: Results of the asymmetry tests described in Section 3.2.1, at 5% nominal size	104
Table 3.3: Results of the asymmetry tests described in Section 3.2.1, at 10% nominal size	105
Table 3.4: Results of the asymmetry tests described in Section 3.2.2, at 5% and 10% nominal sizes	106
Table 3.5: True contingency table under observability of S	107
Table 3.6: Observed tabulation based on an estimator \hat{S} of S	107
Table 3.7: Odds ratio by period and asymmetry test performed. Power = 95%	108
Table 3.8: Odds ratio by period and asymmetry test performed. Power = 80%	109
Table 3.9: Odds ratio by period and asymmetry test performed. Power = 65%	110
Table 3.10: Proportion of tables with odds ratio different from unity	111
Table 3.11: Randomness in the sequence of $\{r_{(i)}\}$ ordered by asymmetry . .	111
Table 3.12: Results of the Z_A test for equality of empirical distributions. . .	112
Table 3.13: Results of the \hat{D} test for equality of density estimates	112
Table 3.14: Results of the Wilcoxon test for location shift	113
Table 3.15: Results of the tests for scale shift	113

ACKNOWLEDGEMENTS

I would like to express a deepest gratitude to Professor Silke Forbes for her guidance, without which this project wouldn't be possible. I would like to thank my committee members, the Professors Karsten Hansen, Ivana Komunjer, Dominique Lauga and David Miller, for their invaluable feedback since earlier stages of this research. I also thank my collaborators Aren Megerdichian and Regio Martins, as well as my friends Marius Rodriguez, Danielken Molina and Philip Neary for their helpful comments and suggestions.

Additionally, I would like to thank my family and friends, who always believed in my potential. Ana deserves special recognition for tolerating my occasional outbursts of frustration. Thank you all for the love and constant support, and thank God for San Diego's year-around pleasant weather.

Chapter 2 is based on *Product Downsizing and Hidden Price Changes in the Ready-to-Eat Cereal Market*, joint work with Aren Megerdichian.

Chapter 3 is based on *Economic Regulation and Asymmetry in Ex-Post Stock Returns*, joint work with Regio Martins.

VITA

- 2001 Bachelor of Science – Economics, State University of São Paulo (USP), Ribeirão Preto, SP, Brazil
- 2003 Master of Arts – Economics, Pontifical Catholic University of Rio de Janeiro, Rio de Janeiro, RJ, Brazil
- 2011 Doctor of Philosophy – Economics University of California, San Diego, CA, United States of America.

ABSTRACT OF THE DISSERTATION

**Essays on Unconventional Pricing Strategies and Impacts of Economic
Regulation on Stock Return Asymmetry**

by

Daniel F. Lima

Doctor of Philosophy in Economics

University of California, San Diego, 2011

Professor Silke Forbes, Chair

This dissertation contains two essays in industrial organization and one related to the corporate finance literature. In Chapter 1, I investigate how consumers' feedback might affect investments in innovation and quality assurance (QA). I also focus on how innovation impacts intertemporal price discrimination. In the model, a monopolist releases a new technology embedded in a durable good. Then, I investigate the determinants of the introduction of a new generation in the second period. In the proposed framework consumers' feedback can reduce the firm's QA/R&D expenses. However, consumers usually provide feedback through complaints, causing reputation damage. I derive conditions under which the lower "quality" version of the good (first generation) will have a higher price than the im-

proved version (second generation). Planned obsolescence causes the price schedule to be steeper than predicted by the previous literature on durable goods. The model predicts and explains Apple's experience with several products, in particular the iPhone family.

In Chapter 2, Aren Megerdichian and I examine a firm's decision to raise price overtly (by increasing the dollar amount of a good) versus the adoption of hidden price change (by decreasing the contents in a good's package). We provide an oligopoly model explaining the hidden price change phenomenon, as well as a comprehensive set of empirical analyses, including demand estimation to assess the impact of hidden price increases on expenditure share and profitability. We focus on the ready-to-eat cereal industry. During July 2007, General Mills decreased the cereal content for 20 out of 23 of their products in our sample of scanner data. We find that some General Mills products gained expenditure share after the hidden price change relative to what the demand system predicts, indicating that a sufficient proportion of consumers did not notice the hidden price change. We also find that some products lost share relative to what the demand system predicts. A key finding is that consumers tend to notice hidden price changes on smaller-sized boxes of cereal, leading them to substitute to larger-sized boxes of cereal.

The final chapter is a joint work with Regio Martins. Our conjecture is that regulated firms may be subject to some regulatory practices that can potentially affect the symmetry of the distribution of their future profits. If these practices are anticipated by investors in the stock market, the pattern of asymmetry in the empirical distribution of stock returns may differ among regulated and non-regulated companies. We review recently proposed asymmetry measures that are robust to the empirical features of return data and investigate whether there are any meaningful differences in the distribution of asymmetry between these two groups of companies.

Chapter 1

Customer Feedback and Planned Obsolescence of High-Tech Products

1.1 Introduction

The past decades witnessed the acceleration of technological progress and, consequently, determining the purchase timing of durable goods has become a more difficult task. Aside from physical obsolescence, the fast pace of technological innovation triggers the speeding of economic obsolescence. In an environment where goods receive frequent upgrades, consumers will be more prone to replace their possessions, even though they are still operational. We can safely assume that producers of high-tech durable goods are well aware of this fact, and that they surely take actions to benefit from it.

Another important aspect about the high-tech durable goods industry is that consumers can actively (yet non-strategically) influence the innovation process. We often observe that after breakthrough technologies/devices and product functionalities are introduced, their refinements are strongly influenced by the experience of their earlier adopters. At the same time, earlier adopters share their experience privately with firms and/or publicly (through the Internet, for example) gives rise to a real world phenomenon: firms, especially those engaged in the production of high-tech durable goods, frequently benefit from consumers' feedback on their product development processes. In this study we focus attention on complaints as a particular channel through which consumers might provide valuable feedback.

The notion presented above contrasts with the paradigm established in the marketing literature. Schibrowsky and Lapidus (1994), for example, claim that “the cost of not doing things right the first time is approximately 20 percent of gross sales for manufacturing companies and 30 percent for service industries [US Congress, 1988]. (...) these findings indicate that the firms that do the best job of reducing customer complaints have a distinct competitive advantage.” I demonstrate that this may not always be the case. Consumers' complaints can be viewed as an important *production factor* available to firms. After all, there might be nothing unintended in Apple's release of a product that “delivers more misses than hits.”¹

¹Article title: “iPhone's list of strengths falls far short of its many disappointing weak-

In order to analyze innovation in durable good markets we need a departure from the traditional literature, otherwise we will eventually reproduce Coase's (1972) monopolist time consistency problem. According to Coase (1972), buyers with rational expectations foresee that the last units sold of a product will be priced at marginal cost; hence, they will be unwilling to pay any price above marginal cost at any point in time.² On the contrary, we show that by timing the innovation to differentiate his product across time the monopolist can recover at least part of his market power.

Lee and Lee (1998) is a good example of departure from Coase's result. The authors introduce a model where price discrimination emerges from changes in consumers' valuation, caused by changes in product quality. In their two-period model, a monopolist who produces a durable good can make a lump sum investment to offer an improved version of the good in the second period. There are two types of consumers, and no considerations about the firm's cost structure. They also assume the existence of an upgrade policy, so the monopolist while marketing the new generation of its product is able to discriminate price not only between consumers with different valuations, but also between consumers with different purchase histories.³

This paper generalizes the analysis in Lee and Lee (1998) by endogenizing the innovation cost and by suppressing the need of an upgrading policy (a salient feature in the software industry, but not as common in the hardware industry). In the model, the actions taken by the monopolist in the first period (producer's choices for price and quality level) determine the amount/relevance of the information that he or she recovers from earlier adopters. Such actions also determine the degree of reputation damage inflicted to the firm by consumers' complaints (reputation damage is not present in Lee and Lee investigation). Differently from

nesses." Published: July 10, 2007. See http://www.infoworld.com/article/07/07/10/28TCiphoneproscons_1.html?source=NLC-PLATFORMS&cgd=2007-07-11.

²Bulow (1982) and Stokey (1981) rigorously establish conditions under which the monopolist may have to give-up monopoly power in equilibrium for particular demand functions. Gul, Sonnenschein, and Wilson (1986) proved for more general demand structures.

³The requisite of an upgrade policy is a non-standard tool in this literature, but fairly employed in the literature on secondhand markets. Tirole (1988) discuss impacts of upgrade policies and the lack of anonymity in a monopolistic market.

their case, I show that serious consideration of consumers' complaints by the firm might increase social welfare.

On the producer's side, consumers' feedback will affect the firm's cost structure in two ways: first, by reducing quality assurance (QA) and research and development (R&D) costs; second, by potentially causing reputation damage due to unseen flaws and/or lack of essential capabilities built-in the first generation. Therefore, the monopolist solves a tradeoff between information acquisition and reputation damage. Adding to it the fact that the monopolist can affect consumers' valuations by timing the introduction of innovation, I am able to model economic obsolescence as an endogenous decision and to derive an optimal price schedule that is steeper than suggested by previous work.⁴

The paper is organized as following. Section 2 discusses some channels by which consumers' feedback can impact on the decisions of a monopolist who intends to release an innovative product. Those channels are well documented (both implicitly and explicitly) by the press and producers' reports/communications. Section 3 introduces the two-period model used to analyze the interactions between consumers and the monopolist with regard to pricing and quality determination. We establish the conditions for an equilibrium where the lower "quality" version of the good (first generation) will have a higher price than the improved version (second generation). Unlike the traditional literature, the monopolist might have to charge a price in the second period that is lower than the lowest valuation, in order to attract the users of the first generation back to the market. Section 4 presents Apple's experience with the iPhone family of products, providing empirical evidence that supports the predictions of the model. Section 5 concludes the paper.

⁴Nair (2007), for example, have studied the video-game industry. The author concludes that the optimal pricing policy should be flatter in the presence of forward-looking consumers in the market. In his specification there is no innovation, and as a consequence consumers who buy the product are not in the market during subsequent periods. However, an endogenous planned obsolescence will attract consumers back to the market, suggesting that the firms can recover at least part of their the market power as time passes, reducing the incentive to lower the price of the first generation.

1.2 Information Acquisition *vs.* Reputation Damage

Investment in innovation and quality assurance is usually an expensive and risky undertaking. For these reasons it is reasonable to presume that producers will not ignore any factors that could reduce both cost and uncertainty related to their business. In spite of that, the analysis conducted by Lee and Lee (1998) reserved no role for cost dynamics. The authors assume that R&D is financed by an exogenous lump-sum cost, and that marginal cost is constant across different generations of the good. Their model setup has two related implications: (i) the opportunity cost of manufacturing the product is constant across time; (ii) the information learned from the demand since the release of the previous generation does not affect the costs of innovating and “bug-fixing.”

This paper attributes a special role for consumers’ feedback in reducing costs and uncertainty faced by the firm⁵. First, because feedback may reduce costs by saving time and resources necessary to the identification of product imperfections. In this sense, feedback can be seen as a form of *strategic outsourcing*, since the firm is getting at least part of its QA task performed by consumers. Strategic outsourcing is a topic that has grown in importance during the last two decades in the management literature, but not very much explored in the economic literature.⁶

QA outsourcing is an established practice in several industry sectors. In particular, software companies have the tradition of publicly recognizing consumers’ participation in the product development process and quality assurance, as evidenced by the following announcement by Microsoft:

Windows Vista Service Pack 1 is an update to Windows Vista that addresses feedback from our customers. In addition to previously released updates, SP1 contains changes focused on addressing specific reliability and performance issues, supporting new types of hardware, and adding

⁵It is important to emphasize that although consumers realize that their feedback is important to further development the product, they are atomic and as a consequence do not behave strategically in the analysis below. As a result, they do not act intending to influence the monopolist’s decisions.

⁶See Quinn (1999).

*support for several emerging standards. Windows Vista SP1 also addresses some management, deployment, and support challenges.*⁷

Such statements are not as common for hardware producers, although we can clearly see that new generations of high-tech equipments do incorporate characteristics or fix imperfections pointed out by earlier adopters. For example, the first generation of Apple's iPod touch does not have a built-in microchip that enables bluetooth services. Because bluetooth headsets and speakers have become very popular (since they work with virtually any cellphone), iPod users started to complain on Apple forums about the lack of connectivity between their new MP3 players/Internet devices and their accessories. Not surprisingly, iPod's second generation is bluetooth enabled. Apple provides yet another example. The iPhone's camera didn't record videos at a time when such capability was common in other devices. Then, after uncountable posts on the Internet about this drawback, the problem was fixed for the next generation.

Additionally, by properly guiding the innovation process consumers' feedback can also reduce the uncertainty faced by firms. If the innovation process is "demand-driven", then consumers' expectations are more likely to be fulfilled, increasing the probability that the product will be successful. As a consequence, the risk of a market failure can be lessened, increasing the firm's expected profitability.⁸

Unfortunately for our monopolist, consumers' feedback is not exclusively related to positive effects. Since complaints are the usual channel through which consumers provide feedback, gathering feedback is frequently associated with brand-devaluing. For this reason, the monopolist has to solve a tradeoff between information acquisition and reputation damage: at the same time that consumers' feedback may reduce costs and uncertainty, complaints can also cause reputation damage. The harm caused to the firm's reputation may negatively affect the sales of future generations due to the fact that consumers might start to distrust the

⁷Published on March 18, 2008. See <http://technet.microsoft.com/en-us/windowsvista/bb738089.aspx?tapm=A80S01A01>.

⁸The magazine article on Forbes *Innovation: A Happy Meal For McDonald's* states that "Consumers provide feedback throughout the process, increasing the odds that McDonald's launches winning products." This statement exemplifies the perception that firms in general explore consumers feedback on their innovation process (see http://www.forbes.com/2007/08/31/christensen-innovation-mcdonalds-pf-guru_in_cc_0904christensen_inl.html).

company. We will investigate the theoretical aspects of this tradeoff in the analysis presented below. To the best of my knowledge, this paper provides the first rigorous and comprehensive treatment of the subject.

The key aspect about the high-tech product markets that we will rely on is that, despite the fact that such products can be purchased by a variety of consumers, some of their built-in functions —specially breakthrough functionalities— are more appealing to *high-skilled* users than they are to “naive” consumers.⁹ This distinction is important because we will consider that sophisticated consumers are more likely to provide valuable information about products than non-sophisticated consumers, helping the firm to foster innovations. We need to keep in mind that since the firm is by assumption introducing a revolutionary product, the release of what we call a “lower quality” version of the product (a concept that will be made more precise later) will create enough enthusiasm among consumers, to the extent that the launching of the product will create sufficient incentive for sophisticated users to engage in earlier adoption.

One question remains: why would sophisticated users be willing to buy the first generation of a product, paying higher prices, if they anticipate the existence of technical glitches and the future release of an improved version of the product? A possible (and plausible) answer could be impatience: sophisticated consumers may be more impatient than other consumers. Alternatively, their behavior could be justified by the fact that high-skilled users might get around problems more easily than other consumers. Products always have a variety of defects, which differ in terms of the fraction of consumers who will be affected by them. Rational consumers would form priors over the probability that they will face problems, and then compare the cost of dealing with problems with the benefits from getting the product sooner at a higher price.¹⁰ It is this latter interpretation that inspires the setup of the model below.

Other conjecture is that *intensive* users (who may not be considered so-

⁹I use the terms high-skilled, advanced and sophisticated interchangeably throughout the paper.

¹⁰I thank Professors David Miller and Roger Gordon from UCSD Economics Department for providing me with this alternative explanation.

phisticated in the sense used above) should be the earlier adopters of products carrying innovation. The justification is that such consumers will benefit from earlier adoption whenever they believe that even a coarser version of the product could help them to perform their personal and/or business activities in a more efficient manner. In this case, intensive users will be the ones thoroughly exploring the product's new functionalities and timely reporting their experiences to the firm.

Now, regarding the supply side, the monopolist may decide to release a not completely developed version of the good to avoid further QA and R&D costs with a product that might not be well accepted by consumers. Additionally, by releasing the product sooner the firm acts preemptively, and might be successful (*i*) in avoiding to be outgone by competitors/developers that may release of similar product in advance, seeking brand visibility; (*ii*) in generating network externalities to preclude potential competitors from releasing a similar product in the subsequent periods; (*iii*) or in promoting brand loyalty, which will assure the incumbent's high market-share in future periods, when competitors enter the market.

1.3 The Model

The specification of the demand expands Lee and Lee (1998) work. Regarding the supply, I depart from previous work by explicitly modeling the impacts of feedback on the firm's innovation and quality assurance costs. Here a monopolist sells a high-tech product to a population of heterogeneous consumers, whose types are indexed by $\theta \in \{L, H\}$. Only consumers know their own types. However, the fraction of the population represented by each consumer type is public information. Sophisticated consumers are indexed by H , and other consumers are indexed by L .

The monopolist sets prices and quality levels in a two-period economy, $t = 1, 2$. The monopolist cannot credibly commit to a pricing policy in which the monopoly price is charged in both periods, and consumers are forward-looking. It is a well established result that consumers with perfect foresight have incentives to

postpone their purchases. Along those lines, Nair (2007) concludes that the optimal pricing policy for durable goods is not steep. In the equilibrium analysis below I will keep the perfect foresight assumption, while showing that intertemporal price discrimination is still important, perhaps for different reasons than the ones pointed out in the existing literature. Perfect foresight and perfect disclosure of quality are simplifying assumptions necessary to rule out uncertainty from the model.

1.3.1 The Demand

Consider a continuum of consumers lying on the interval $[0, 1 + \mu]$. The size of type L population is normalized to 1, and the size of type H population is $\mu \geq 0$.

The durable good produced in period t provides consumers of type θ with a utility level v_t^θ . As mentioned before, consumers differ in their capabilities to overcome difficulties caused by the product's solvable flaws or incompatibilities. This intuition allows us to write $v_t^H \geq v_t^L$, i.e., sophisticated consumers have *higher valuations* than unsophisticated consumers.

Each consumer uses at most one unit of the durable good per period and there is no discounting for either consumers and the monopolist. It implies that a consumer who purchases in the first period may continue to use the good in the second period, thus obtaining the utility level of the previous period again. If a new product is purchased in the second period, then the old version is scrapped. Additionally, the consumer will purchase the product whenever he or she is indifferent between buying it, or not.¹¹

I introduce in the model Daughety and Reinganum (2008) concept of product quality. The functional form (1.1) implies that all product characteristics that affect a consumer's utility (other than prices) are summarize by an index \tilde{q}_t , for $t = 1, 2$.¹² I further assume that the quality index q_t is observable, thus writing

¹¹We assume that the monopolist is always able to sell all units produced, i.e., the market clearing condition is always satisfied.

¹²Technically, the index \tilde{q}_t is a mapping from a multidimensional quality space into a unidimensional space. It is worth noticing that the number of product's capabilities cannot be considered as a *proxy* for quality, since an increase in the number of functionalities can result in either an improvement or a deterioration of product's overall quality. By simply adding extra capabilities

$\tilde{q}_t = q_t$.¹³ The consumer's utility is then given by

$$v_t^\theta = q_t\alpha + (1 - q_t)(\alpha - \delta_\theta) = \alpha - (1 - q_t)\delta_\theta, \quad (1.1)$$

where $q_t \in (0, 1]$ can be interpreted as the consumers opinion about the product. If $q_t = 1$ consumers view the product as being flawless. For $q_t < 1$ the product is regarded as imperfect, and the lower q_t is, the “more defective” the product is perceived by consumers. I rule out from the analysis the case $q_t = 0$, since it represents a “undesirable” good in the assumed quality space and the monopolist does not have incentive to offer it. The parameter $\alpha > 0$ measures the utility level provided by the flawless version of the good, while δ_θ measures type's θ disutility, for $\theta = L, H$, in the case of a defective unit.

Consumers' heterogeneity is represented by $\delta_\theta \in \{\delta_H, \delta_L\}$. From (1.1) it is easy to see that, *ceteris paribus*, consumer's utility decreases by δ , the cost of dealing with flaws. I assume that $\delta_L \geq \delta_H$, i.e., low type consumers face a higher cost of dealing with flaws.¹⁴ As a result, sophisticated consumers will be more willing to adopt the new technology earlier than naive consumers, since sophisticated consumers forgo relatively more utility when they do not buy the product. In our context, this specification has yet another sensible implication: low type consumers benefit more from quality improvements than high type consumers. Mathematically, $v_2^L - v_1^L \geq v_2^H - v_1^H$ (assuming that the product does not worsen in time, $q_2 \geq q_1$). Intuitively, since type H consumers are less affected by product imperfections, the same quality improvement, $q_2 - q_1 > 0$, translates into lower benefits for them, then it does for unskilled consumers. Moreover, since in my setup (i) demand is increasing in quality, (ii) reputation damage is decreasing in quality and (iii) the marginal cost to produce any $q_2 \in (0, q_1]$ is the same, the monopolist has no incentive to produce a second generation of the good such that

to the product the firm (or the user) may impair functions already implemented and properly working. For example, the addition of several *add-ons* to an Internet browser usually results in internal routine conflicts, which will eventually cause a software freezing. So, although new capabilities are added or cumbersome routines have been automated, the freezing of the software will cause a decrease in its overall satisfactoriness, resulting in a quality decrease in my framework.

¹³In principle, \tilde{q}_t is an unobservable variable to consumers prior to purchase. So the real assumption is that the producer truthfully reveals its value beforehand (perfect disclosure).

¹⁴The intuition is simple: more skillful consumers are less impacted by product defects.

$q_2 \leq q_1$, which implies that $v_2^\theta \geq v_1^\theta$ for $\theta = H, L$.¹⁵

Given prices p_1 and p_2 , for periods $t = 1$ and $t = 2$, respectively, and the assumption of no discounting, consumer surplus for different purchase patterns are presented in Table 1.1. To illustrate how those entries should be read, consider type H possible holdings. For example, the purchase pattern $H.1$ happens when the high-skilled consumer purchases the good in the first period and continues to use it during the second period. If a purchase occurs in period two, then purchase patterns $H.2$, $H.3$ and $H.4$ are the candidates. The difference is that in $H.2$ the consumer delays her purchase of the first generation until the second period, while in $H.3$ the consumer buys only the second generation of the product. $H.4$ happens when she buys both generations. Analogous interpretations hold for the bottom panel of Table 1.1, which displays payoffs for type L consumers.

Consumers' will choose one of the purchase patterns described in Table 1.1, or opt for an outside option. Valuations in Table 1.1 are normalized, such that the consumer surplus provided by the outside option is zero.

1.3.2 The Supply

I now turn attention to the monopolist's problem, where consumer's feedback will play a central role in the firm's cost dynamic. The firm has to solve a tradeoff between information acquisition (which reduces the firm's QA/R&D cost) and the incurrence of reputation damage.

In my setup the firm actively employs intertemporal price discrimination to control the amount and quality of the information collected from consumers. By raising the price charged for the first generation, the firm restrict the number of units sold, and consequently the volume of feedback received from earlier adopters (referred here as the "feedback benefit"). At the same time, higher prices also reduce the reputation damage inflicted to the firm (referred here as "feedback cost"). In this scenario, information acquisition and reputation damage are both increasing functions of first period sales. The intuition is that feedback lowers the firm's

¹⁵This is not necessarily a general result: firms might be able to produce goods of better quality and decide (strategically) to produce lower quality goods. For example, firms might not be willing to release a new technology in the first period if they cannot fully capture the benefits.

QA/R&D cost through the information channel, while potentially revealing to the general public bad news about the product (hurting the firm's prestige by ruining consumer confidence). Naturally, the reputation damage is also a decreasing function of the quality level of the first generation, since a better product results in less complaints. The cost function below captures such features:

$$C(Q_1, Q_2, q_1, q_2) = \gamma q_1^2 + \chi_1 Q_1 + \frac{\gamma(\max\{q_2 - q_1, 0\})^2}{1 + f(Q_1)} + \frac{\phi Q_1}{q_1} + \chi_2 Q_2$$

where $Q_t \in \{\mu, 1, 1 + \mu\}$, $t = 1, 2$, is the quantity sold in each period, and as before, q_1 and q_2 are the quality levels of the first and second periods, respectively. The parameters χ_1, χ_2, γ and ϕ are positive. The factor $\gamma/(1 + f(Q_1))$ captures the information intensity of the feedback benefit. The firm's R&D and QA own expenditures will solely determine q_1 , while consumers' feedback plus firm's further R&D and QA investment will determine the level of q_2 .

The QA and R&D cost components are (i) γq_1^2 , associated with the investment made by the firm previous to the product release, (ii) and $\gamma(\max\{q_2 - q_1, 0\})^2/(1 + f(Q_1))$, associated with additional QA/R&D investments necessary for the introduction of an improved version of the product. As functions of q , both QA/R&D cost components are convex, meaning that additional QA/R&D investments to obtain quality improvements occur at increasing rates. The $\max\{\cdot\}$ operator captures the notion that the monopolist is able to produce at any quality level up to q_1 at no additional cost during the second period.

In the denominator of the last ratio we have the function $f(Q_1)$, which assumes the value $\mu\beta_H$ if only type H buys the product during $t = 1$, and the value $\mu\beta_H + \beta_L$ if both consumer types buy the product in $t = 1$. I am essentially using the quantity sold in $t = 1$, Q_1 , as a *proxy* for the amount of information recovered from the earlier adopters. The division by $f(Q_1)$, satisfying $\Delta f(Q_1) > 0$ for $\Delta Q_1 > 0$, accounts for the fact that consumers' feedback foster innovation, cheapening costs and reducing the uncertainty related to the innovation process (by promoting a demand-driven R&D). Feedback also reduce the cost of assuring optimal quality standards (QA outsourcing). This information channel effect lowers the monopolist's cost in the second period.

The parameters β_H and β_L determine the relative information “density” of type H and type L feedback. According to the discussion above, I assume that $\mu\beta_H \geq \beta_L$. This assumption expresses the fact that the amount of information transmitted through feedback increases at a decreasing rate. In other words, the “information content” of type H complaints is more relevant/useful to the monopolist than the information gained from type L complaints.

Reputation damage is quantified by the term $\phi Q_1/q_1$. Low investments in QA/R&D during the first period implies a lower q_1 , resulting in more reputation damage. Q_1 is also used as a *proxy* for the volume of complaints, and consequently for the amplitude of the bad news spreading. A broader dissemination of bad news will result in greater harm caused to the firm. The intuition is simple: a million people merely complaining about iPhone’s inability to handle Flash and Java components is more harmful to Apple than the complaint of a single IT guy from Silicon Valley, who may even alert Apple to the existence of *open source* solutions for the problem. In general, “technology geeks” communicate technical issues directly to producers or discuss them in specialized forums, instead of using the mass media as a communication channel.¹⁶

A monopolist facing the cost structure described above solves the following problem:¹⁷

$$\max_{(p_1, p_2, q_1, q_2)} \pi = p_1 Q_1 + p_2 Q_2 - \gamma q_1^2 - \chi_1 Q_1 - \frac{\gamma(\max\{q_2 - q_1, 0\})^2}{1 + f(Q_1)} - \frac{\phi Q_1}{q_1} - \chi_2 Q_2,$$

subject to previously stated restrictions on model parameters and the following participation and incentive-compatibility constraints,

1. $V_H(Q_H, P_H) - V_H(Q_L, P_L) \geq 0$
2. $V_L(Q_L, P_L) - V_L(Q_H, P_H) \geq 0$

¹⁶In fact, managers know that mass media personnel regularly inspect those forums and could disseminate the bad news about their products, increasing the harm on reputation. Still, the media does not have access to complaints that have arrived directly to the firm and media personnel is not aware of the existence of several forums. Furthermore, some general interest publications (such as BusinessWeek) have a technology page where a “tech guy” reviews a product and gives opinions about the product’s features and shortcomings. But, even in this case, the specialist usually discuss ways to improve the product.

¹⁷Remembering that the discount rate is set to zero.

$$3. V_H(Q_H, P_H) - V_H(0, 0) \geq 0$$

$$4. V_L(Q_L, P_L) - V_L(0, 0) \geq 0$$

where Q_t, q_t and $p_t, t = 1, 2$, and cost parameters are defined as above. V_H is type H 's two-period indirect utility function, and V_L is type L 's two-period indirect utility function. Q_H, Q_L, P_H and P_L are type H and type L two-period purchase histories and associated price vectors, respectively.

1.3.3 Equilibrium Analysis and Comparative Statics

Now that I have described the demand and supply setups, I can study the monopolist's strategies. A theoretical result worth mentioning is that a monopolist selling a durable good can not credibly commit to a pricing strategy, which involves no price cuts over time, due to time-inconsistency associated with this policy.¹⁸ The game played between consumers and the monopolist has multiple equilibria, however I will restrict attention to a particular pure-strategies symmetric (in the sense that consumers of the same type play the same pure strategy) subgame perfect Nash equilibrium. Below I determine and interpret the conditions necessary for the existence of the "planned economic obsolescence equilibrium."

For the sake of tractability, I assume the following simplifying assumptions. First, without loss of generality due to the discrete nature of the demand, let's assume that the producer can only choose between two levels of quality, therefore $q_t \in \{\underline{q}, \bar{q}\}, 0 < \underline{q} \leq \bar{q} \leq 1$. A discrete q has the following impact on consumers' valuations: when the monopolist advertises $q = \underline{q}$, valuations are given by \underline{v} , analogously the consumer's valuation is denoted \bar{v} when $q = \bar{q}$.

Second, in order to concentrate attention on the impacts of feedback on the cost structure, I will further assume that $\chi_t = 0$ for $t = 1, 2$. However, there are reasons to believe that $\chi_1 > \chi_2$ (learning-by-doing achievements, for example). Readers are encouraged to reincorporate χ_1 and χ_2 into the analysis to see that $\chi_1 > \chi_2$ would relax the restrictions imposed on parameters to support the equilibrium that I am investigating.

¹⁸After the first period the monopolist has incentive to deviate from fixed price strategy. See Stokey (1979) and Stokey (1981).

By incorporating the assumptions above I can rewrite the objective function as

$$\max_{(p_1, p_2, q_1, q_2)} \pi = p_1 Q_1 + p_2 Q_2 - \gamma \left(q_1^2 + \frac{(q_2 - q_1)^2}{f(Q_1)} \right) - \frac{\phi Q_1}{q_1}. \quad (1.2)$$

Following Lee and Lee (1998), I impose restriction on the parameters in order “to focus on interesting cases in the spirit of the classical durable goods monopoly problem.” Basically, such restrictions guarantee that market power matters, i.e., that the monopolist finds optimal to restrict the output. The first restriction assures that if there is no innovation, then the monopolist will set a price that will make the product attractive only to consumers of type H in period one. The second restriction ensures that if both types buy the product in the first period, then the monopolist will benefit from selling the updated version only to the type H during period two. Those restrictions are

$$(i) \quad v_1^H > v_1^L(1 + \mu)/\mu$$

$$(ii) \quad v_2^H > v_2^L(1 + \mu)/\mu$$

The expressions in (i) and (ii) are slightly different from their counterparts in Lee and Lee (1998) for two reasons.¹⁹ First, I do not assume the necessity of an upgrading policy to derive the results. The expression analogous to (ii) in their paper is a statement about the optimality of their upgrade pricing, while my restriction must hold in general, independently of consumers purchase histories. Second, they use the population sizes switched: μ is the size of the type L population in their paper, while it is the size of the type H population in this paper. I decided to relabel the population sizes to emphasize that the monopolist’s decision is particularly affected by the ratio of high-type consumers in it. Now we are ready to establish the main result in the paper, stated in Proposition 1 below.

Proposition 1.3.1 *The consideration of the effects of consumers’ feedback on innovation and reputation damage enable a new sort of equilibrium, one that allows*

¹⁹Their analogous expressions are

$$(i') \quad v_1^H > (1 + \mu)v_1^L$$

$$(ii') \quad v_2^H - v_1^H > (1 + \mu)(v_2^L - v_1^L).$$

the following strategies to configure a sub-game perfect Nash equilibrium (SPNE) without requiring the existence of an upgrading policy.

- monopolist sets $\{p_1 = \underline{v}^H, q_1 = \underline{q}\}$ and $\{p_2 \leq \bar{v}^L, q_2 = \bar{q}\}$. Previous work do not indicate the possibility of $p_2 < \bar{v}^L$.
- type L don't buy in $t = 1$, but buys in $t = 2$ since $p_2 \leq \bar{v}^L$.
- type H buys in both periods.

Proof See Appendix A. Table 1.2 summarizes the conditions supporting Proposition 2.3.1.

The next step in my investigation is to provide interpretations for the conditions supporting the studied equilibrium. An important implication is that we do not need “too many” type H consumers to generate the appropriated incentives for intertemporal price and quality discrimination. The conditions (L.2) and (M.1) will hold if $\mu < 1$.

A higher β_H relaxes all conditions where it appears. The intuition is simple: the better type H consumers are on providing feedback, the greater is the incentive for the monopolist to outsource QA and to rely on their feedback to drive the innovation process. The direct welfare implication is that consumers' feedback boost technology advance in our setup, since the range of p_2 that supports (A.1) is increasing in μ and β .

Although I establish a positive association between feedback and the amount sold during previous period, the useful “information density” obtained through the feedback provided by sophisticated consumers is enough to reduce costs while balancing against reputation damage. Low type consumers are in larger number in the market and even if they can provide the same level of information density than high types, they will cause a bigger reputation damage by spreading bad news more widely.

Also, as expected an increases of μ would relax conditions (M.1) – (M.4). By relaxing (M.1) and (M.2) we benefit the monopolist in two ways: first, she has more consumers to sell the product to and, consequently, to extract surpluses

in period one. Second, there are more consumers to be brought back to the market in period two via quality improvement. The impact of an increase of μ on (M.5) is ambiguous. (M.5) represents the tradeoff between the amortization of the QA/R&D cost and the reputation damage.

Note that an increase in the quality differential will not necessarily relax the restrictions in Table (1.2), favoring the selected equilibrium. Either a higher level for \bar{q} or a lower \underline{q} will not necessarily relax (M.4) or (M.5), but it will relax (M.3). In addition:

Proposition 1.3.2 *(L.1) and (L.2) together imply (L.3).*

Proof *The monopolist optimally fix $p_1 = \underline{v}^H$. As a consequence (L.3) becomes $p_2 \leq \bar{v}^L + \underline{v}^H - 2\underline{v}^L$. But according to (L.2) $\underline{v}^H - 2\underline{v}^L > 0$ and the result follows.*

■

Given Proposition 2, the most interesting possibility happens when (L.1) holds strictly. In this case a new sort of result emerges. The price charged by the monopolist is lower than the valuation of the low type consumer in $t = 2$. The intuition is that the monopolist has to fix p_2 low enough to attract high-type consumers back to the market, and if such price is sufficiently low, type L consumers will get positive consumer surplus, an extra incentive for naive consumers to delay purchase.

Finally, the higher the cost parameter γ , the more attractive the set of strategies presented in Proposition 1 is to the monopolist. The intuition behind this result is straightforward: if QA/R&D cost increases, then the firm will benefit even more from the consumers' feedback, since returns for innovation becomes less risky when innovation is demand-driven and quality assurance becomes cheaper when the monopolist has her consumers replacing the in-house team on a diligent search for product imperfections. On the other hand, ϕ (the parameter that represents reputation damage) has exactly the opposite impact. A higher ϕ makes less likely that the firm will select the SPNE analyzed in this paper among the multiple equilibria for this game.

1.4 Case Study: Apples's iPhone Release Strategy

Empirical evidence supporting the SPNE studied in the previous section is provided by Figure 1.1. It shows that the general public delayed their purchase until the release of an improved and cheaper version of the product. The information about iPhone sales is compiled from Apple's Quarterly Financial Result Reports available at the company's website.²⁰ We see that as time passes new generations are introduced, there is a sharp increase in iPhones's sales. In fact, sales increase tremendously after the release of second generation in July 11, 2008. The first generation total sales was 6.1 millions units, over a one-year period, approximately, while the second generation sales reached 6.9 millions units during the quarter after its introduction.²¹

The iPhone's first generation cost the consumer \$599 for the 8-gigabyte model and \$499 for the 4-gigabyte model, on a two-year contract with AT&T in the USA market. There were plentiful allegations on the media claiming that such prices were too high for a tie-in sale of a device and a two-year service contract.

Apples's release strategy has certainly separated high-valuation consumers from other potential buyers, reducing demand for the first generation, and consequently alleviating problems caused by frustration among consumers (the main source of reputation damage). At the same time, as suggested by the model, earlier adopters abundantly shared their experiences with Apple. A quick search on the Internet promptly proves that consumers and specialists have heavily complained about the first generation's missing and flaws.²² Complaints provided the guidance for the development of the second generation, which addressed most of the alleged issues. In fact, Apple has a specific webpage to collect customer feedback about

²⁰The link to the online version of the figure is http://upload.wikimedia.org/wikipedia/commons/8/8f/IPhone_sales_per_quarter.svg and links to Apple's financial statements can be found at http://en.wikipedia.org/wiki/File:IPhone_sales_per_quarter.svg#Data_and_references.

²¹The complete timeline of Apple products can be accessed at http://en.wikipedia.org/wiki/Timeline_of_Apple_products#2000s.

²²See footnote 1 for reference, for example.

the iPhone.²³

The second generation arrives in the market costing much less and addressing most part of the consumers complaints. It reaches the market costing \$199 for an improved 8-gigabyte device that only 12 months ago cost \$599, and \$299 for a 16-gigabyte version.²⁴

An important question remains: has Apple been successful in attracting type *H* consumers back to the market during subsequent releases? The answer is yes. According to a survey by Piper Jaffray of early iPhone adopters, 38% of the second generation buyers owned the first generation, while 56% of the earlier adopters of the third generation upgrade from a previous iPhone model. A recent article on the Forbes magazine concludes that 77% of iPhone 4 sales were upgrades. Our model successfully explains Apple's experience with the iPhone family of products.²⁵

Other company benefiting from consumers' feedback and practicing a pricing policy consistent with the results of the previous section is Amazon with their e-book reader, Kindle. Despite the fact that Amazon does not release official sales information about its e-book, reports by market analysts suggest that Amazon's Kindle experience is very similar to Apple's experience with the iPod family (which is very similar to the iPhone experience). The Business Insider stated that "initial skepticism about Amazon's Kindle is being replaced by euphoria: Citi's Mark Mahaney, who was already bullish on the e-book reader, declares that is indeed going to be Amazon's iPod."

Amazon introduced the Kindle in November 2007, selling it for \$399. The company had the device sold out in five and a half hours after release, and it remained out of stock for the next five months, until late April 2008. Such practice is a clear signal that the firm wanted to restrict the sales of the first generation.²⁶

²³The following statement "we read all feedback carefully" can be found at <http://www.apple.com/feedback/iphone.html>.

²⁴An 8-gigabyte maximum memory capacity was one of the major complaints at that time, along with longer battery life and the need of a faster processor for a better Internet experience.

²⁵<http://tech.fortune.cnn.com/2010/06/25/77-of-iphone-4-sales-were-upgrades/>.

²⁶"Amazon was out of stock of the Kindles through Christmas and right into the new year. Consumers who placed an order anyway were told their Kindle would be delivered in 10-13 weeks. For those consumers who were willing to wait, they got a nice surprise. They were shipped

On February 23, 2009, the Kindle 2 became available for purchase for the price of \$359 (reduced to \$299 short after, on July 8, 2009).²⁷

1.5 Conclusion

The first generation of high-tech products are often shipped with a variety of imperfections and incompatibilities, which impact different types of consumers in different ways. In general, high-skilled users get around problems more easily than other consumers, suggesting that they are less affected by product flaws (equivalently, product flaws are less costly for those consumers). As a result, sophisticated users might have higher valuations for such products than the average consumer, becoming stronger candidates for earlier adoption.

The view adopted in this paper—that a producer of high-tech durable goods can benefit from consumers’ complaints—contrasts with the paradigm established in the marketing literature. This paper demonstrates theoretically that consumers’ complaints can sometimes be considered as an important production factor. It also provides empirical evidence that high-tech companies play the strategies described in Proposition 2.3.1.

The main result can be summarized as follows. The firm recognizes the high-skilled consumers incentive to earlier adoption and sets a very high price for the first generation of the product. By doing so, the firm restricts sales to a smaller fraction of the population. According with assumptions, consumers more capable of delivering valuable feedback and less prone to spread bad news to the general public will be the ones buying the first generation of the good. Thereby, the firm is able to limit reputation damage and at the same time collect useful information about the product. Using this information, the firm implements a demand-driven innovation process. Finally, the monopolist sells the second generation to all con-

the brand new Amazon Kindle 2 generation as it hit the Amazon website in February 2009!” <http://ebookreadersresource.com/ebook-reader-reviews/amazon-kindle-2-review/>.

²⁷During the last Amazon.com shareholder meeting (May 28, 2009) Amazon’s Chief Executive Jeff Bezos states that “the company may never release sales figures for the Kindle: there’s a competitive advantage in keeping the numbers close. You may just have to remain curious,” Bezos said in response to an audience member’s question.

sumers. Product improvements and a low price for the new generation induce previous period consumers to upgrade their possessions as well.

In a previous treatment of economic obsolescence, Lee and Lee (1998) concluded that the cannibalization of the stream of rents before technological innovation is not desirable. Contrasting with their result, in this paper such cannibalization is desirable. The main difference between the two papers is that cost dynamics are ignored in their analysis, while I model how consumers' feedback can directly and indirectly affect it. The tradeoff between information acquisition and reputation damage is a key factor determining how the strategies described in Proposition 1 emerge as a SPNE. Another relevant difference between the two works is that I preserve consumer's anonymity, while Lee and Lee (1998) assume the existence of an upgrade policy allowing the monopolist to track consumers purchase histories (a feature not usually observed in the hardware industry).

Interestingly, the existence of the equilibrium analyzed does not require a large fraction of the population to be represented by high-valuation consumers. Proper incentives materialize even when the population of high-type consumers to be brought back to the market in the second period is smaller than the population of low-type consumers. More striking is the fact that in order to induce high-types to upgrade, the monopolist might need to fix a price for the second generation that is lower than the valuation attributed by type L , leaving a positive surplus to the low valuation consumer. This result implies a steeper pricing schedule than previous work and contrasts with traditional results of demand screening.

In this paper I assume that the monopolist keeps its monopoly power during both periods. Future work will introduce competition and analyze the more realistic situation, in which there is a monopoly during the first period and that monopolist becomes an incumbent firm in the second period.²⁸

As mentioned before, the model presented in this paper predicts and explains Apple's experience with the iPhone product family. We verify that (i) consumers provide feedback; (ii) the company collects feedback and implements

²⁸This situation is usually verified after breakthrough innovations. For example, at present Apple's iPhone faces arduous competition from Android phones (phones using Google's operational system).

the capabilities suggested and fixes the flaws pointed by consumers; *(iii)* the company releases an improved and cheaper version of the product; *(iv)* sales skyrocket in the second period; and *(v)* least, but not last important, consumers do upgrade their possessions. Whether Apple’s described experience is a particular case, or the “birth of a new paradigm” for the production of technology intensive products is an empirical question that needs more observation, and a more thorough treatment. Regardless, to the best of my knowledge this paper is the first to rigorously model how consumers’s feedback may affect the firm cost dynamic and influence the firm’s decisions.

1.6 Tables and Figures




Figure 1.1: iPhone Sales Chart

Table 1.1: Consumer Surplus

Purchase Pattern	Holdings		Consumer Surplus
	1st period	2nd period	
	1st generation	2nd generation	
Panel A: Type H Consumers			
$H.1$	✓	✓	$2v_1^H - p_1 = 2[\alpha - (1 - q_1)\delta_H] - p_1$
$H.2$	-	✓	$v_1^H - p_2 = \alpha - (1 - q_1)\delta_H - p_2$
$H.3$	-	-	$v_2^H - p_2 = \alpha - (1 - q_2)\delta_H - p_2$
$H.4$	✓	-	$v_1^H + v_2^H - p_1 - p_2 = 2\alpha - (2 - q_1 - q_2)\delta_H - p_1 - p_2$
Panel B: Type L Consumers			
$L.1$	✓	✓	$2v_1^L - p_1 = 2[\alpha - (1 - q_1)\delta_L] - p_1$
$L.2$	-	✓	$v_1^L - p_2 = \alpha - (1 - q_1)\delta_L - p_2$
$L.3$	-	-	$v_2^L - p_2 = \alpha - (1 - q_2)\delta_L - p_2$
$L.4$	✓	-	$v_1^L + v_2^L - p_1 - p_2 = 2\alpha - (2 - q_1 - q_2)\delta_L - p_1 - p_2$

Table 1.2: Conditions Supporting Proposition 1 (grouped by players)**Type H:**

$$H.1) p_2 \leq \bar{v}^H - \underline{v}^H$$

$$H.2) \underline{v}^H - p_1 \geq 0$$

$$H.3) \Delta v^H \leq \Delta v^L$$

Type L:

$$L.1) p_2 \leq \bar{v}^L$$

$$L.2) \underline{v}^L < \underline{v}^H/2$$

$$L.3) p_2 \leq \bar{v}^L + p_1 - 2\underline{v}^L$$

Monopolist:

$$M.1) \underline{v}^H > \underline{v}^L(1 + \mu)/\mu$$

$$M.2) \bar{v}^H > \bar{v}^L(1 + \mu)/\mu$$

$$M.3) \underline{v}^L < \gamma \underline{q}^2 + \phi \mu / \underline{q}$$

$$M.4) p_2 \geq \underline{v}^L / (1 + \mu) + \gamma (\bar{q} - \underline{q})^2 / (1 + \mu)(1 + \mu \beta_H)$$

$$M.5) p_2 \geq [(2\bar{v}^H - \underline{v}^H)\mu + \bar{v}^L + \phi \mu (\bar{q} - \underline{q}) / \bar{q} \underline{q} - \gamma [2(\bar{q} \underline{q} - \underline{q}^2) + \mu \beta_H (\bar{q}^2 - \underline{q}^2)] / (1 + \mu \beta_H)] / (1 + \mu)$$

Chapter 2

Product Downsizing and Hidden Price Changes in the Ready-to-Eat Cereal Market

2.1 Introduction

Firms have at least three ways to change the prices of their products. The first way, defined here as an overt price change, is to change the dollar amount charged per package without changing the amount of the good contained in the package. The second way to change price, defined here as a hidden price change, is to change the package size without changing the dollar amount charged for the package. Finally, a third possible way to change price is with a combination of an overt and hidden price change, defined here as a hybrid price change.

There are very few prior studies on hidden price changes. The only two relevant studies we are aware of are Adams, di Benedetto, and Chandran (1991) and Gourville and Koehler (2004), neither of which directly examines the impact on a firm's revenues or profits arising from a hidden price change. We therefore contribute to the existing literature by studying the extent to which a firm is successful in increasing its products' prices with a hidden price change. If a sufficient number of consumers do not respond to the hidden price increase, then we expect the firm to benefit from the price increase. In contrast, if a sufficient number of consumers do notice the hidden price increase and thus substitute to other goods, then we expect the firm to suffer from the price increase.

Furthermore, we examine hidden price changes in the ready-to-eat breakfast cereal industry. During 2006 and 2007, the prices of commodities that are primary inputs for cereal production, such as wheat, rice, and corn, increased dramatically. This increase in input costs prompted General Mills to increase prices for nearly all of its products with a hidden (or hybrid) price change in July 2007. For example, General Mills changed its flagship line of Cheerios products from 10oz., 15oz., and 20oz. down to 8.9oz., 14oz., and 18oz., respectively. We employ an extremely rich scanner data set that allows us to determine which products had a price change and how the price change was implemented. The detail of the data also allows us to undertake a variety of empirical analyses, including estimating a full demand system to examine the success of a hidden price change, which no other study has done to date.

We conduct four empirical investigations to determine the impact of hidden

price changes on General Mills' cereal business. First, we estimate a basic linear demand model as in Gourville and Koehler (2004), and we also improve their specification by including prices of all cereals and promotional variables. Second, we estimate demand with the almost ideal demand system (AIDS) to predict expenditure share for the after-downsizing sample period, and compare it to the actual expenditure share. We also estimate a difference-in-difference specification for expenditure share to determine how General Mills' portfolio of products that were downsized fared as a whole. Finally, we examine Megerdichian's (2009) analysis of product downsizing's effect on General Mills' profitability.

We find that out of the 20 General Mills products in our sample that downsized cereal content, about half the products increased expenditure share by more than what the AIDS predicts, indicating a sufficient number of consumers didn't notice the product downsizing for these goods. A key finding is that the products that did benefit from the downsizing tend to be large size boxes (by weight), while those that did not benefit from the downsizing tend to be small size boxes. One explanation is that consumers are more likely to notice hidden price changes when already small boxes become even smaller, leading them to substitute to larger boxes.

With a difference-in-difference estimation scheme, we find that the downsizing had a negligible effect on the expenditure share for the 20 General Mills products as a whole. While this result may appear to suggest that consumers noticed the hidden price change, it is perhaps more plausible the result just confirms the AIDS prediction exercise in that some of General Mills' products benefitted because enough consumer didn't notice, while other products did not because consumers did notice, thus leading to a negligible net effect.

This paper is organized as follows. Section 2 provides an overview of the hidden price change phenomenon, including a review of prior relevant studies on the topic. Section 3 introduces an extended version of the AIDS model and a leadership game used to derive testable implications about consumers reaction to hidden price changes. Section 4 describes the supermarket scanner data on cereals that we employ here. In Section 5, we present a variety of empirical work outlined

above. Section 6 concludes.

2.2 Hidden Price Changes

Firms have more than one way to change price. The first way, defined here as an overt price change, is to change the dollar amount charged per package without changing the amount of content in the package. For example, a cereal manufacturer may increase price from \$4.00 per 16oz. box to \$4.50 per 16oz. box, a 12.5% price increase. The second way to change price, defined here as a hidden price change, is to change the amount of content in a package without changing the dollar amount charged for the package. For example, a cereal manufacturer may decrease the size of a 16oz. box to 14.22oz. per box while keeping the price fixed at \$4.00 per box, also constituting a 12.5% price increase. Consistent with previous literature on this topic, we refer to the act of decreasing the amount of content in a package as *downsizing*. Finally, a firm may opt for some combination of an overt and hidden price change, defined here as a hybrid price change. For example, a firm can decrease the size and price of a \$4.00 per 16oz. box to \$3.75 per 12.69oz. box, which is again a 12.5 % increase in price measured in dollars per ounce.

There are important aspects of a hidden price change worth noting. First, a hidden price change is tantamount to an overt price change, but in a way that some consumers may not notice (hence, hidden). Second, changing the amount of content in a package may require redesigning the package to accommodate the new size, or at the very least require the firm to change the weight information on the package. Thus, a hidden price change may involve production considerations, whereas an overt price change does not. Finally, it may be relatively difficult to undo a hidden price change by going back to the pre-downsized package; thus, undoing a hidden price change may require a counteracting overt price change. We examine these ideas in subsequent sections.

Gourville and Koehler (2004) refer to several product segments that have undergone downsizing, including coffee, yogurt, potato chips, baby diapers, bot-

tled water, breakfast cereal, and ice cream. Adams, di Benedetto, and Chandran (1991) similarly point to these industries, and also note several other product segments, including candy bars and paper towels. Notably, the examples have the commonality of being packaged goods that sell primarily in supermarkets and, for the most part, appear to be differentiated products in an oligopoly-type market. That the primary outlet for these products are supermarkets is an important point, as temporary price promotions in supermarkets account for a significant part of a good's price variation (see Pesendorfer, 2002, and Hosken and Reiffen, 2004).¹ Therefore, downsizing may be more likely to go unnoticed when sticker prices vary considerably week-to-week.

2.2.1 Cost as a Driver of Downsizing

A reason manufacturers may change the price of a product is to remain profit maximizing when faced with a change in marginal cost. Consider a profit maximizing monopolist that sets price according to the well-known Lerner price-cost margin

$$\frac{p - c}{p} = -\frac{1}{\eta},$$

where η is the own price elasticity of demand, p is price, and c is marginal cost. Rearranging gives the familiar optimal pricing rule

$$p = \left(\frac{\eta}{1 + \eta} \right) c.$$

According to this pricing equation, a firm may find it necessary to change price in response to a change in its marginal cost. During the last several years leading up to the start of the 2008 U.S. recession, prices of input commodities, including petroleum, grains, and rice increased markedly.

Figure 2.1 presents monthly price indices for several commodities that are inputs for the production of breakfast cereals. Prices of corn, rice, wheat, and petroleum were on the rise between 2003 and 2007, and increased considerably starting in 2006.² The price of sugar, on the other hand, was relatively stable over

¹See also Figure 2.5 that conveys the volatility of retail prices for General Mills Cheerios.

²Petroleum can be thought of as a proxy for general energy costs incurred to produce and transport cereal.

the five-year time period. Table 2.1 presents yearly data for the commodities. Corn and wheat had large price movements in 2006 and 2007, while rice and petroleum had their largest price movements in 2004 and 2005. According to Megerdichian (2009), 60% of the cereal products in his sample list corn or wheat as the primary ingredient. It is clear that cereal manufacturers faced rising input prices during the last decade, particularly during 2006 and 2007, which drove some of them to react with a hidden price increase; specifically, General Mills downsized its boxes for 20 out of its 23 products in our sample in July 2007.

Moreover, as one industry commentator reported in December 2007, “to counter rising commodity expenses, General Mills has reduced merchandise price discounts and promotions... General Mills even shrank its cereal-box size and increased the average price of cereal per ounce to help counter rising costs and better compete against Kellogg Co., its larger rival in cereals... These actions, on top of price increases for select items and new product introductions, are expected to offset this cost pressure...”³

If the variable input cost and thus marginal cost of production for a firm increases, the firm will find it necessary to increase the price of its products to remain profit maximizing. The question of interest is should the manufacturer overtly increase the price of the product or should it change price through a hidden increase? The answer to this question depends in large part on how consumers react to overt versus hidden price changes.

Consider the example given earlier of changing the price of a 16oz. box of cereal from \$4.00 to \$4.50 (a 12.5% increase) versus attaining an equivalent price increase by keeping the price at \$4.00, but downsizing the box from 16oz. to 14.22oz. Assume that the intermediary retailer simply passes on price changes to the consumer. If consumers consider the two different pricing schemes to be completely equivalent —i.e., they are rational— then it does not matter which way firms increase price since consumers will perceive both identical. If enough consumers do not perceive them to be equivalent, potentially because consumers simply don’t notice that the contents of the box changed from 16oz. to 14.22oz.,

³“General Mills Backs Year View as Profit Rises,” by Matt Andrejczak and Angela Moore, www.marketwatch.com, December 19, 2007.

then the manufacturer should engage in a hidden price increase.

Of course, a manufacturer can't keep decreasing the contents of its package, because at some point it becomes obvious to the consumer that they aren't getting much for their money. In fact, our empirical estimates show that, on average, larger-sized cereal products gained share relative to smaller-sized products after the downsizing.

2.2.2 Previous Research on Downsizing

To date, the most comprehensive academic study examining product downsizing and the hidden price change phenomenon is Gourville and Koehler (2004). Although they refer to several product segments that engage in hidden price changes, an indication that it is a fairly common occurrence, there are very few academic studies that rigorously examine product downsizing. The few studies we found on the topic are tangentially relevant, as none of them adequately measure the quantitative impact of hidden price changes.

Granger and Billson (1972) conduct an experiment with consumers to determine the extent to which they are aware of the price per unit of measurement (e.g., dollar per ounce) of soft drinks and laundry detergents. They find that consumers are generally not aware of the price per unit. When the price per unit of measurement is explicitly revealed to the consumers, they tend to substitute towards larger-sized items that typically have a lower price per unit. Note, however, that this study took place prior to retailers adopting unit pricing standards. As of 2008, 19 states in the U.S. have adopted unit pricing regulations that requires retailers to present price per unit of measurement information;⁴ the majority of retailers, particularly supermarkets, voluntarily offer such information.⁵ From this study we conclude that consumers tend to be more sensitive to overt price changes than hidden price changes.

Adams, di Benedetto, and Chandran (1991) offer several reasons for why firms find the need to downsize packaging, including maintaining a price point, in-

⁴National Inst. of Standards and Technology, Uniform Laws and Regulations, NIST Handbook 130, 2009.

⁵Food Marketing Institute, Item and Unit Pricing, October 2007 (electronic article).

creasing margin and profitability, increasing frequency of purchase, and offsetting raw material cost increases. They identify 25 product segments that committed downsizing primarily during the 1980s, including candy bars, coffee, cereal, ketchup, paper towels, and soap. In many instances, they document that the same packaging is used to hold the downsized contents.

Their study, although informative, provides very little empirical evidence assessing whether firms are successful in implementing hidden price increases. For example, they examine three brands of cereal that were downsized during the early 1980s, and make inferences about the impact on expenditures after the downsize by simply examining expenditures after the downsize. However, many other confounding changes may take place during the relevant time period after the downsize that would lead to incorrect inferences. For example, if some of the many other cereals in the market that are substitutes for the downsized product decreased their price, or if promotions for the downsized product increases after the hidden price change, this may lead to an increase in expenditures for the downsized product for reasons potentially unrelated to the hidden price increase. We address these concerns in our empirical study in Section 5.

Finally, Gupta et al. (2006) examine the legal and ethical issues related to product downsizing and hidden price increases. They maintain that although hidden price changes through product downsizing may adversely affect consumers that do not notice the price increase in the same way as they would notice if it were an overt price increase, manufacturers are nevertheless in compliance with U.S. Food and Drug Administration regulations. However, they point out that there may be room for the Federal Trade Commission to pursue firms if the hidden price change is viewed as deceptive. To date, we are not aware of any legal or regulatory consequences resulting from hidden price changes. The reason is likely due to the fact that most retailers, regardless of which state they are located, provide consumers with a dollar per unit of measurement price.⁶

⁶For example, most grocery stores present a shelf sticker that states a 16oz. box of cereal is \$3.99; the same sticker will state \$0.25/oz. below the \$3.99 figure, although typically in much smaller font.

Gourville and Koehler (2004)

We pay special attention to Gourville and Koehler's (2004) Harvard Business School working paper (hereafter, GK), as it is the most complete study that directly examines product downsizing and hidden price changes. GK examine multiple product categories, using both surveys and market data, to determine that consumers react to overt price increases much more strongly than to hidden price increases. GK conduct four studies to assess hidden price increases; we describe and assess three of the relevant ones here.

First, GK examine the package size and shelf price of ready-to-eat cereals from a single supermarket at one point in time. They suggest that if consumers are more sensitive to overt price changes relative to hidden price changes, and if manufacturers are aware of this, then we should expect to see more variation in package sizes than in prices. For 157 cereal products, they examine the ratio of the standard deviation of package sizes to the mean of package sizes, and compare this to a similar ratio for price. For example, they find that the ratio for package size is 0.244 for General Mills products compared to a ratio for price of 0.122 for the same General Mills products. They conclude that because there is more variation in package sizes relative to prices, manufacturers are exploiting the fact that consumers don't notice hidden price changes relative to overt price changes.

We don't necessarily dispute GK's conclusion that consumers are less sensitive to hidden price changes, but the way they implemented their empirical test is not appropriate. Assuming that examining variance is applicable here, then what makes sense is to examine the variance of shelf price changes for a good over time versus the variance of its package size changes over time. In other words, their idea to examine variance may be valid, but it should be done as a time-series exercise for each good separately, rather than a cross-sectional exercise for all goods at one point in time.

For example, let s_t denote a good's package size in time t , $d_t \equiv 1[\text{downsize}]$, $\delta \equiv s_b - s_a$, where s_b and s_a are the sizes of the good before and after the downsizing, respectively. Consider a relevant time period $t = 1, 2, \dots, T$, and that the product downsizes once during that period. The mean and the standard deviation of s_t is

given by $\bar{s} = T^{-1} \sum_t s_t = (1 - \pi)s_b + \pi s_a = s_b - \pi\delta$ and $\sigma = \delta\sqrt{\pi(1 - \pi)}$, respectively, where $\pi \equiv T^{-1} \sum_t d_t$. If the product downsizes once, $\sqrt{\pi(1 - \pi)} \in (0, .5]$. Following GK's idea to use σ/\bar{s} to evaluate the impact of product downsizing, we have that $\sigma/\bar{s} = \delta/(\delta + 2s_a)$ for the extreme case, where $\sqrt{\pi(1 - \pi)} = .5$. It does not take much for σ/\bar{s} to be very small. Consider a good that is downsized from 20 ounces to 18 ounces exactly at the halfway point of a relevant time period that lasts $T = 260$ weeks. In this case, $\delta = 2$, $\bar{s} = 20$, $\sigma = 1$, and GK's ratio is $\sigma/\bar{s} = 0.05$, a very small number when examining downsizing correctly as a time series exercise, which would in this case indicate that downsizing is not undertaken very often by the firm, but does not tell us much about how profitable the hidden price change was.

The fact that GK find variation in package sizes for a cross section of products at a particular point in time is not surprising for a differentiated products market. It is well known that at any point in time cereals come in a wide variety of shapes, sizes, and densities, among other attributes, to satisfy consumers' preferences, which includes preferences for different package sizes. GK's analysis is simply picking up the fact that, for example, Cheerios comes in three different packages sizes.

In their second study, GK survey 60 adults about how they would price coffee if they were in charge of pricing at a gourmet coffee shop. Their survey asks subjects whether it is better to price coffee at \$6 per half pound or \$12 per pound; the subjects are also asked to indicate which of the pricing schemes would promote sales better for the store.⁷ Nearly 87 percent of respondents chose the \$6 per half pound pricing option, and the mean score is 2.85 for the second question, indicating that respondents believe the \$6 per half pound pricing scheme to be more lucrative for the business even though the two pricing schemes are the same price in terms of dollars per pound. Respondents justify their answers by noting that consumers won't notice the "per half pound" part of the price, or that the unit of measurement (the denominator of price) is not as important to consumers

⁷The latter question is based on a one to seven scale. A score of one is described to be that \$6 per half pound is "much more effective," while a score of seven is described to be that \$12 per pound is "much more effective," and four is described to be that both are "equally effective."

than the dollars (the numerator of price). GK find that this sentiment favors a hidden price increase strategy for firms.

In their fourth study, GK return to the cereal industry, but this time use data covering 145 weeks during the late 1990s for an undisclosed number of supermarkets that they apparently have aggregated into a single representative store. Their data is for four cereal products from a single manufacturer that engaged in product downsizing during the relevant time period, and contains information on unit sales, price, and package size. They specify a standard linear regression equation with quantity purchases of cereal (in boxes) as a function of a trend variable, own price (presumably, in dollars per box), and size (ounces per box).

GK explain that if consumers do not notice hidden price changes, then the price variable will have more impact on quantity than the size variable in their regressions. For their sample of four cereals, they obtain statistically significant negative coefficients on price that are relatively large in magnitude compared to their estimates of the size coefficient. They conclude that this is evidence that consumers pay more attention to overt price changes than hidden price changes from downsizing. We note here that although it's not fair to compare the coefficient estimates on price and size directly, because price and size have different units of measurement, comparing the respective t-statistics of price and size, which are unit free, still shows price dominating size in their regressions. However, it might only reflect the fact the the price variable has more variability than size.

Except for one of the four cereals, the coefficient on size is not statistically significant, and GK note that the negative signs on the size variable are “counter-intuitive.” However, there is an intuitive explanation for the negative coefficients. Letting q denote quantity purchases and s size, their estimate for three of the four products is $\partial q/\partial s < 0$, meaning that a decrease in the amount of cereal in a package leads to an increase in the number of boxes of cereal purchased. The reason this is actually intuitive is that we may expect that, *ceteris paribus*, consumers who don't notice the downsizing and who consume a fixed amount of cereal during a relevant time period to purchase more boxes of cereal to maintain their consumption if there is less content in each box. The demand model we present

bellow is flexible to accommodate this possibility.

There are a few empirical issues with their analysis. First, as GK themselves correctly point out, they have left out marketing-mix variables, such as local advertising and in-store displays, that typically play an important role in driving purchases in supermarkets. Second, they are essentially estimating a poorly-specified linear demand equation for each of the four cereals in their sample by leaving out the prices of the many other cereals in the market.⁸ Third, it appears the price variable they use is the regular shelf price for the cereal, not the effective transaction price that takes into account temporary price promotions. The effective transaction price for cereals during a promotional week can sometimes be half of the shelf price, which may impact their regression results (see Figure 2.5). Lastly, it appears they have taken data from multiple supermarkets and aggregated it into a single representative supermarket that they use for their regressions. This may effect the results, particularly the size coefficient, due to aggregation bias.

We do not disagree with GK's overall conclusion that consumers are less sensitive to hidden price changes versus overt price changes. Nevertheless, none of their studies adequately addresses whether the hidden price change was successful from the firm's point of view. They only address consumers' perceptions about hidden price changes, and then draw conclusions about how those perceptions may affect firms that engage in product downsizing. In contrast, we examine the impact of hidden price changes on a firm's revenue and profitability, and we fix most of the problems with their quantitative demand analysis.

2.3 The Model

Deaton and Muellbauer's (1980) almost ideal demand system (AIDS) is rooted in consumer theory, and is considered to have a high level of econometric flexibility, in that even if the true demand function that describes the data is not the AIDS, the AIDS model will provide an approximate estimate that is near the true DGP. Moreover, the AIDS specification has expenditure share as the dependent

⁸They do not intentionally leave out prices of other cereals (potentially substitutes) and marketing-mix variables; it appears they simply don't have the data.

variable, and is thus a convenient framework to determine how successful, from a revenue share perspective, GM’s product downsizing strategy was in 2007. In this section we extend the AIDS model to incorporate the possibility that consumers may fail to recognize implicit price changes. We present this extended version of the model to help the interpretation of the empirical results below. Subsequently, we analyze under which circumstances firms can benefit from this “bounded rationality.” Our approach is simple: first we aggregate the two kinds of consumers separately, then we combine their demand functions to examine the equilibrium in a Nash-Bertrand duopoly model with leadership.⁹

2.3.1 Extending the AIDS Model

We can interpret the market demand function in DM’s work as an aggregation over consumers that are always able to perfectly recognize price changes, even when prices are adjusted implicitly (as defined above). In contrast, the model introduced here allows us to distinguish between explicit and implicit price changes by assuming the existence of two types of consumers. There is a proportion ρ of the population that is unable to identify *hidden* price changes, and a proportion $1 - \rho$ that is fully rational.

In the model the inability to identify implicit price changes is product-specific, and the probability of failing to recognize an implicit price change is denoted by $\rho_j, j = 1, 2, \dots, J$. In particular, we believe that ρ_j is a function of the box size and brand loyalty. We assume that the effect of brand loyalty on ρ_j is positive, i.e., people will fail to notice hidden price changes more often when they are *ex-ante* less willing to switch brands (what causes them to pay less attention to price tags). Regarding the box size, our testable conjecture is that consumers are more likely to recognize implicit price changes in smaller boxes. The intuition is that once the box size is sufficiently small/light, a further downsizing becomes easily recognizable, rendering hidden price changes less effective.

A market $m = 1, 2, \dots, M$, is defined as a city-retailer-week combination;

⁹We have chosen to analyze the first-mover advantage because the observed characteristics of the read-to-eat cereal market resembles a leadership competition.

i.e., the index m subsumes r (city-retailer) and t (week). The demand is a function of the product's price per pound in a given market, p_{jm} , and box size is used as a control variable in our regressions. The idea is that consumers' choice can be at least partially explained by the amount of product that they want to stock, in order to have enough units lasting until their next trip to the grocery store.

The most relevant distinction between implicit and explicit price changes is that whenever the latter happens, all consumers update their information set to incorporate the new price per pound, whereas when the former happens a fraction ρ of consumers is going to fail to update it. For this reason, price *decreases* are always implemented explicitly. It is in the producers' best interest that all consumers unmistakably recognize a price decrease, since there is no competitive advantage in masquerading a price decline. As a result, implicit price change is equivalent to downsizing and we use these terms interchangeably.

Let's define the downsizing indicator variable $d \in D = \{0, 1\}$. It assumes the value zero if there is no change in content, and one otherwise. The *extended-demand* in market m for product j is given by

$$w_{jm} = \begin{cases} \lambda_j + \sum_k \gamma_{jk} \ln p_{km} + \beta_j \ln \left(\frac{X_m}{P_m} \right) + \text{controls} & \text{if } d = 0 \quad (2.1a) \\ \tilde{\lambda}_j + (1 - \rho_j) \left(\sum_k \gamma_{jk} \ln p_{km} + \beta_j \ln \left(\frac{X_m}{P_m} \right) \right) + \\ \rho_j \left(\sum_k \gamma_{jk} \ln p'_{km} + \beta_j \ln \left(\frac{X'_m}{P'_m} \right) \right) + \text{controls} & \text{if } d = 1, \quad (2.1b) \end{cases}$$

where w_{jm} is the expenditure share based on dollar sales of good j in market m , X_m is the total expenditure of all $k = 1, 2, \dots, K$ cereal products in market m , and P_m is the Stone Price Index, given by $\ln P_m = \sum_k w_{km} \ln p_{km}$. The variables X'_m and P'_m are defined analogously, but calculated using p'_m , the price vector in consumers' information set previously to the price change. The *controls* used are discussed in empirical section below.¹⁰ λ , γ , β and ρ are parameters.

¹⁰Other possibility would be to measure the market share in terms of the number of boxes sold

The expected signs for the coefficients are as follow: positive for the cross-price effects γ_{jk} , necessary for positive cross-price elasticities, and negative for the own price effects γ_{jj} , in order to get negative own-price elasticity. Betas are either positive or negative.

Our demand specification incorporates the traditional AIDS model as a special case when $d = 0$ and/or $\rho_j = 0$. Additionally, when $d = 1$ and $\rho \neq 0$, equation (2.1b) allows prices to be incorrectly observed (p'_m) by a fraction ρ of consumers. The main problem is that p'_{jm} is not observed by the firm or by the econometrician, and needs to be estimated or *proxied* for.

In the present analysis we assume that p'_m is the previous period price vector, thus downsizing implies that $p_{jm} - p'_{jm} \geq 0$. This price misperception makes possible that producers increasing prices implicitly experience an attenuated negative impact on expenditure shares (when compared to the impact produced by an explicit price change). Moreover, in an environment where all other firms are increasing prices explicitly, a successfully implemented downsizing might even result in an increase in the firm's expenditure shares.

Despite the fact that firms can benefit from the price effect mentioned in the previous paragraph, downsizing is not a dominant strategy because it also affects λ_j . Note that when $d = 1$,

$$\tilde{\lambda}_j = \rho_j \lambda'_j + (1 - \rho_j) \lambda_j. \quad (2.2)$$

Two opposite forces will determine if $\tilde{\lambda}_j$ is greater or less than λ_j . Gupta et al. (2007) state that “package downsizing (...) has the potential to mislead customers in the buying process due to an unfavorable balance of information within the dyad. This could give rise to serious moral and ethical consideration.” Our model incorporates this suggested possibility by the authors. Consumers may prefer explicit price changes and penalize firms for practicing downsizing. If “reputation damage” is the prevailing effect, then $\lambda'_j < \lambda_j \Rightarrow \tilde{\lambda}_j < \lambda_j$, affecting product's j expenditure share negatively.

On the other hand, downsizing can also induce new substitutions patterns,

 (i.e. quantities), but the Almost Ideal Demand System model is derived from consumers' expenditure functions, consequently implying market shares measured in terms of relative revenue.

mainly between goods belonging to the same family of products, i.e., products that differ only in terms of box sizes. For example, assume that there is a specific type of cereal that is sold in two different box sizes, and that the producer decides to downsize both boxes. All consumers who are unable to recognize the downsizing would still buy the same boxes in the same amount, to the extent that for them nothing has been altered.

The reaction of consumers able to recognize the downsizing is more complex. They can switch to another brand or they can switch to a different box size of the same product. If consumers of a particular cereal type have a “strong” preference for it, then we expect that some consumers who have chosen to buy the *old* smaller box (the smaller box carrying the content pre-downsizing) will purchase the *new* bigger box, which can be viewed as more similar to the old smaller box by those consumers. At the same time, we expect some consumers who were buying the bigger box before the downsizing to continue to purchase the larger size. As a result, λ'_j can be greater than λ_j for some products, although most likely not for all goods belonging to a certain family of products at the same time.

Figure 2.6 illustrates the effects described in the above paragraph. In the picture we consider two brands with two products each. The products of each brand differs only in terms of their box sizes. The prices displayed reflect the fact that bigger boxes (indexed by B) are usually sold at a lower price per pound than smaller boxes (indexed by S). For simplicity, we assume that brand 2 keeps prices per pound constant, and that brand 1 downsizes both products. Brand 1 prices per pound increases from p_1^B to z_1^B and from p_1^S to z_1^S . The arrows represent possible substitution patterns emerging from downsizing. There are two arrows departing from p_1^S and only one arriving at z_1^S , indicating that the quantity of smaller boxes sold by brand 1 can decrease. At the same time, there are three arrows point at z_1^B and two leaving from p_1^B . This represents the possibility that brand 1 can actually increase the number of larger boxes sold after downsizing. Now we move to the next task modeling the supply side to determine under which circumstances firms can benefit from consumers’ price misperception.

2.3.2 The Supply

Liu (2005) introduces a Stackelberg model to analyze the benefits of leadership in a homogeneous good duopoly under demand uncertainty. Inspired by Liu's model, we investigate the impact of demand uncertainty in a differentiated goods market. In order to accomplish this goal, we present a multi-product differentiated Nash-Bertrand duopoly model with leadership. Another remarkable difference with respect to Liu's work is that here uncertainty does not vanish after the leader's move, since demand uncertainty is product specific.

In our study the demand uncertainty arises from the fact that firms do not observe λ' in equation (2.2), i.e., they are unable to determine how consumers will react to implicit price changes beforehand. For tractability, we assume the existence of two firms, a leader (L) and a follower (F), each producing one type of cereal packaged in two different box sizes. All four products are sold in a single market. The game has two stages; marginal and fixed costs are set to zero, and there is no barriers to entry.¹¹ Quantities are measured in pounds and prices in dollars per pound.

Firm L has to decide whether or not to downsize *both* products during the first stage of the game.¹² The leader sets prices p_L^B and p_L^S for the bigger and smaller boxes, respectively. In order to avoid ambiguity, prices are denoted by z_L^B and z_L^S whenever downsizing happens. Firm F then sets prices explicitly in the second stage of the game. From the firm's perspective, the parameter λ' in equation (2.2) has uniform distribution on $[\underline{\lambda'}, \overline{\lambda'}]$.

¹¹In the ready-to-eat cereal industry there are two natural candidates for the market leadership. As presented later, Kellogg and General Mills market shares place them as leader competitors in the United States. Yet, there is no evidence that one of them should be considered as leader regardless the time period. Given the co-movement of their market shares, we can assume that half of the time each firm would be playing the role of leader. In our sample period General Mills played the downsizing strategy first, and for this reason we can interpret the model implications with GM as the first mover and other producers grouped as a single follower.

¹²We restrict our analysis to the downsizing of all products of a given producer because that is the pattern observed in our dataset.

2.3.3 Partial Equilibrium Analysis

Given L 's pre-committed prices, F solves the following problem

$$\max_{\{p_F^B, p_F^S\}} \Pi_F(p_L^B, p_L^S) = Q_F^B p_F^B + Q_F^S p_F^S.$$

Dividing the profit function above by the total revenue, we can rewrite F 's problem in terms of equation (2.1a), as follows

$$\max_{\{p_F^B, p_F^S\}} \pi_F(p_L^B, p_L^S) = w_F^B + w_F^S. \quad (2.3)$$

We assume the existence of an equilibrium in pure strategies and a strictly positive support for prices, thus the first-order conditions imply that

$$p_F^B = -\frac{2\gamma_F^B}{A(\beta_F^B + \beta_F^S)} \quad (2.4a)$$

$$\text{and } p_F^S = -\frac{2\gamma_F^S}{B(\beta_F^B + \beta_F^S)}.^{13} \quad (2.4b)$$

In the benchmark case the leader sets prices overtly as well. The leader problem's can be written as $\max_{p_L^B, p_L^S} \pi_L = (w_L^B + w_L^S | p_F^B, p_F^S)$, where w_L^B and w_L^S are defined as in equation (2.1a). and the conditional reaction functions p_F^B and p_F^S are the follower's optimal pricing policies, given by (2.4a) and (2.4b), respectively. It follows that

$$p_L^B = -\frac{2\gamma_L^B}{2\Gamma_L^B + C(\beta_L^B + \beta_L^S)} \quad (2.5a)$$

$$\text{and } p_L^S = -\frac{2\gamma_L^S}{2\Gamma_L^S + D(\beta_L^B + \beta_L^S)}.^{14} \quad (2.5b)$$

¹³where

$$A \equiv \sum_{\substack{j=F,L \\ i=B,S}} \left(\frac{\partial Q_j^i}{\partial p_F^B} \frac{p_j^i}{X_m} - \frac{\partial w_j^i}{\partial p_F^B} \ln p_j^i \right); B \equiv \sum_{\substack{j=F,L \\ i=B,S}} \left(\frac{\partial Q_j^i}{\partial p_F^S} \frac{p_j^i}{X_m} - \frac{\partial w_j^i}{\partial p_F^S} \ln p_j^i \right).$$

¹⁴where

$$\Gamma_L^B \equiv \frac{\gamma_F^B}{p_F^B} \frac{\partial p_F^B}{\partial p_L^B} + \frac{\gamma_F^S}{p_F^S} \frac{\partial p_F^S}{\partial p_L^B}; C \equiv \sum_{\substack{j=F,L \\ i=B,S}} \left(\frac{\partial Q_j^i}{\partial p_L^B} \frac{p_j^i}{X_m} - \frac{\partial w_j^i}{\partial p_L^B} \ln p_j^i \right)$$

In order to determine the leader's optimal prices under downsizing, we define the price vectors $z \equiv (z_L^B, z_L^S, p_F^B, p_F^S)$ and $p \equiv (p_L^B, p_L^S, p_F^B, p_F^S)$. When implementing the downsizing strategy, firm L maximizes the following expected profit problem:

$$\max_{z_L^B, z_L^S} E_\lambda[w_L^B + w_L^S | p_F^B, p_F^S],$$

where w_L^B and w_L^S are now defined as in equation (2.1b).

In a two-stage model with one market $p'_j \equiv p_{j,t-1}$, for all j , thus the partial derivative of any function of p'_j with respect to $p_{j,t}$ is zero.¹⁵ Optimal prices are given by

$$z_L^B = -\frac{(2 - \rho_L^B - \rho_L^S)\gamma_L^B}{(2 - \rho_L^B - \rho_L^S)\Gamma_L^B(z) + C(z)((1 - \rho_L^B)\beta_L^B + (1 - \rho_L^S)\beta_L^S)} \quad (2.6a)$$

$$\text{and } z_L^S = -\frac{(2 - \rho_L^B - \rho_L^S)\gamma_L^S}{(2 - \rho_L^B - \rho_L^S)\Gamma_L^S(z) + D(z)((1 - \rho_L^B)\beta_L^B + (1 - \rho_L^S)\beta_L^S)}, \quad (2.6b)$$

where $\Gamma_L^B(z), \Gamma_L^S(z), C(z)$ and $D(z)$ are defined as $\Gamma_L^B, \Gamma_L^S, C$ and D , respectively, but evaluated at price-vector z , instead of p .

Lemma 2.3.1 *If functions Γ, C and D are nowhere constant, then the prices set explicitly or by downsizing are the same, provided that $\rho_L^B = \rho_L^S \equiv \rho$.*

Proof *After setting $\rho_L^B = \rho_L^S = \rho$ in equations (2.6a) and (2.6b) we can factor them out and compare the resulting expressions to (2.5a) and (2.5b), respectively. If functions Γ, C and D are injective functions we must have $p = z$.*

This result indicates that the leader has incentive to exploit consumers' misperceptions only when product characteristics affect consumers ability to recognize implicit price changes. The ratios ρ_F^B and ρ_F^S do not affect the incentive to downsize because the follower is changing prices overtly. Lemma 2.3.1 allows us to derive the fundamental relationship between optimal prices under overt and hidden price change strategies:

$$\Gamma_L^S \equiv \frac{\gamma_F^B}{p_F^B} \frac{\partial p_F^B}{\partial p_L^S} + \frac{\gamma_F^S}{p_F^S} \frac{\partial p_F^S}{\partial p_L^S}; D \equiv \sum_{\substack{j=F,L \\ i=B,S}} \left(\frac{\partial Q_j^i}{\partial p_L^S} \frac{p_j^i}{X_m} - \frac{\partial w_j^i}{\partial p_L^S} \ln p_j^i \right)$$

¹⁵We assume that consumers observed all prices the first time that they purchased one of the products.

Proposition 2.3.1 *If consumers are more likely to miss the downsizing of bigger boxes, then downsizing will result in higher prices for both products, provided that sensitivity to total real expenditure is positive and stronger for bigger boxes.*

Proof From Lemma 2.3.1 we know that implicit and explicit price changes lead to the same prices whenever $\rho_L^B = \rho_L^S$. Taking the natural logarithm of equations (2.6a) and (2.6b) and calculating the total differentials $d \ln z_L^B$ and $d \ln z_L^S$ with respect to ρ_L^B and ρ_L^S , for small changes $d\rho_L^B = -d\rho_L^S \equiv \varepsilon > 0$, we have

$$d \ln z_L^B = \varepsilon \left(\frac{\partial \ln z_L^B}{\partial \rho_L^B} - \frac{\partial \ln z_L^B}{\partial \rho_L^S} \right) \Big|_p = C(\beta_L^B - \beta_L^S)\varepsilon$$

$$d \ln z_L^S = \varepsilon \left(\frac{\partial \ln z_L^S}{\partial \rho_L^B} - \frac{\partial \ln z_L^S}{\partial \rho_L^S} \right) \Big|_p = D(\beta_L^B - \beta_L^S)\varepsilon$$

Since prices must be positive and Γ is negative (considering that $\partial p_F^j / \partial p_L^k > 0$, for $(i, j) \in \{B, S\}$, in a Bertrand-Nash equilibrium), the following restrictions must apply: (i) $C(\beta_L^B + \beta_L^S) > 0$ and (ii) $D(\beta_L^B + \beta_L^S) > 0$, which imply that $\text{sign}(C) = \text{sign}(D)$. Consequently, $\text{sign}(d \ln z_L^B) = \text{sign}(d \ln z_L^S) = 1$ if $\beta_L^B \geq |\beta_L^S| > 0$.

We interpret Proposition 2.3.1 as follow: first, the leader has incentive to downsize the bigger boxes setting a price per pound above (2.5a) because a higher proportion of consumers will miss implicit price changes for that product, given that $\rho_L^B > \rho_L^S$.¹⁶ Second, after downsizing the “new” big box size might be more comparable to the “old” smaller size. The price increase above (2.5b) induces a higher fraction of consumers to substitute away from the smaller to the newer bigger size box (in addition, some consumers will fail to recognize that price of the bigger box has gone up). In summary, the result in Proposition 2.3.1 is consistent with an increase of the expenditure share for the bigger boxes, at the cost of a decrease of smaller boxes expenditure shares, as our empirical results show.

The last task in this section is to show that even when $\rho_L^B = \rho_L^S$, downsizing is not equivalent to an overt price change, despite the fact that $p = z$. The

¹⁶Perhaps because smaller boxes are already too “light”, triggering consumers to observe downsizing more easily.

uncertainty about λ' causes implicit price change not to be a dominant strategy. Note that when the leader changes prices explicitly, the intercept of the profit function is

$$\lambda_L^B + \lambda_L^S. \quad (2.7)$$

However, under downsizing the intercept of the expected profit function is

$$(1 - \rho_L^B)\lambda_L^B + (1 - \rho_L^S)\lambda_L^S + \frac{\bar{\lambda} - \underline{\lambda}}{2}(\rho_L^B + \rho_L^S). \quad (2.8)$$

The firm may not be indifferent between the two price change policies simply because the intercept of the profit functions are likely to be different. In particular, if $\rho_L^B = \rho_L^S \equiv \rho$, then (2.8) becomes

$$(1 - \rho)(\lambda_L^B + \lambda_L^S) + (\bar{\lambda} - \underline{\lambda})\rho. \quad (2.9)$$

Considering other things equal and the previous discussion, if downsizing causes some degree of reputation damage, then (2.9) is less than (2.7) and implicit price change is a dominated strategy (i.e., if $\lambda_L^B + \lambda_L^S > \bar{\lambda} - \underline{\lambda}$). In this case, the leader forgoes his first mover privilege and opt for an explicit price change. Now that we have determined that hidden price change is not a dominant strategy, in the subsequent sections we present several approaches to empirically investigate if GM benefitted from downsizing in 2007. We start with a brief overview of the ready-to-eat cereal industry and our dataset.

2.4 Data

Currently, the four major nationally-branded cereal manufacturers operating in the U.S. are General Mills, Kellogg, Post, and Quaker. A fifth important source of cereal sales are store brand products, which are typically inexpensive versions of nationally-branded cereal products. We refer the reader to numerous studies by industrial organization and marketing researchers for a background on the cereal market.¹⁷ Table 2.2 presents aggregate market share and pricing data.

¹⁷See, for example, Schmalensee (1978), Ippolito and Mathios (1990), Hausman (1997), Rubinfeld (2000), Nevo (2001), Nevo and Wolfram (2002), Shum (2004), and Megerdichian (2009, 2010).

General Mills and Kellogg are the two largest players in this market, followed by Post, Quaker, and store brands, respectively.

We employ a rich set of supermarket scanner data from Information Resources, Inc. (IRI). The data set consists of variables measuring price, quantity, and promotional variables for 150 of the top-selling cereal products sold in the U.S. for three years between January 2005 and December 2007. The data set has a panel structure, where the time dimension has a weekly frequency and the individual dimension is a city-supermarket chain.¹⁸ There are 157 weeks of data for each of the 121 city-supermarket chains covering 41 cities¹⁹ across the United States. We refer the reader to Megerdichian (2009, 2010) for further details.

Figure 2.2 displays the monthly market share based on pounds of cereal sold of the four cereal manufacturers as well as store brand products. The store brand market share does not deviate much from five percent on a month-to-month basis, while the branded manufacturers have considerably more volatility over time.²⁰ Figure 2.2 indicates that not only do the market shares of Kellogg and General Mills vary over time, but that they are nearly mirror images. The correlation of the monthly market share for Kellogg and General Mills is -0.78 , indicating that Kellogg's gain in market share is typically General Mills' loss, and vice-versa.

Figure 2.3 presents the weighted average monthly retail price per pound for the manufacturers between 2005 and 2007. Figure 2.4 presents the weighted average price per box. General Mills products tend have higher prices than the other manufacturers, while the store brands products are sold at a deep discount compared to the nationally-branded products. What is notable here is that about July 2007, the point during which General Mills downsized 20 of the 23 products in the sample, the price per pound (Figure 2.3) of General Mills products increases, while their price per box (Figure 2.4) during the same time period is relatively flat. This is indicative of their hidden price increase campaign. Moreover, the timing of the increase in the price per pound for General Mills is roughly in line with the

¹⁸The data are not for individual stores; rather, they report information for a supermarket chain, which may be comprised of many stores in a given city or district of a city.

¹⁹A city is approximately equivalent to the Census Bureau's metropolitan statistical area (MSA) or combined metropolitan statistical area (CMSA).

²⁰For more details, see Megerdichian (2009).

sky-rocketing commodity prices presented in Figure 2.1.

The prices variation over time is due to temporary promotions that are more pronounced in the disaggregated weekly data. Figure 2.5 presents the weekly price per pound and price per box of General Mills Cheerios 15oz. (14oz. after the downsizing), the best selling product in the sample over the three year period. The series is for a single supermarket chain located in a major U.S. city.²¹ The numerous price decreases are the result of temporary promotions, which are an important source of price variation in this data. The richness of this data set permits us to determine when downsizing took place for General Mills products in 2007. We explore this next.

2.4.1 Identifying Downsizing

Importantly, the disaggregated nature of the data allows us to identify when General Mills (GM) downsized their products.²² GM downsized most of its cereal products in 2007, and downsized the remainder in 2008. Kellogg started downsizing its products in 2008. To note, 2008 and 2009 are beyond our sample.

Dividing a good's price in dollars per box by its corresponding price in dollars per pound yields the ratio pounds per box, or ounces per box after multiplying by 16. This ratio is constant over time for products that did not downsize, but decreases if the product underwent a downsize. Figure 2.5 presents the ounces per box ratio for GM Cheerios 15oz. using the price series in the same figure. The product downsizing for this good took place at this particular supermarket chain in early July 2007, as GM downsized its product from 15 ounces to 14 ounces. This can be seen in two ways. First, the gap between the two price series increases during July 2007. Second, the ounces per box, calculated as the ratio of the two price series, drops from 15oz. to 14oz. during July 2007.

Tables 2.3 and 2.4 present the 20 GM products that underwent downsizing

²¹Confidentiality agreements entered into by Megerdichian for his UCSD PhD dissertation do not permit him to reveal the identities of the supermarkets.

²²General Mills downsized a substantial portion of its cereal products in 2007, and downsized the remainder in 2008. Kellogg started downsizing its products in 2008. To note, 2008 and 2009 are beyond our sample.

during July 2007.²³ Because different supermarkets changed over to the new, smaller boxes at different times, the downsizing date reported in Table 2.3 is an average across all city-supermarket chains in the sample. The variable ounces per box reports the amount of cereal in the box before and after the downsizing. For example, GM changed the 10oz. box of Cheerios to a 8.9oz. box near mid-July 2007. This is an 11% decrease in the amount of cereal in the box, which translates into an implied hidden price increase of 12.4%.

Tables 2.3 and 2.4 also provide average prices for two measures: price per pound and price per box. Consider again the Cheerios product that changed from 10 ounces to 8.9 ounces. The average price per box across all the city-supermarket-weeks is \$2.95 before the downsizing, and is \$2.96 after the downsizing. Yet the average price per pound jumps 11.9%, from \$4.73 per pound to \$5.29 per pound. This is a clear-cut example of a pure hidden price change, which can also be seen by comparing the implied price change with the change in the price per pound, which is about the same for the 10oz. Cheerios product.

The same table shows that some products, on average, underwent a hybrid price change, whereby the product had both downsizing and an overt price change. Nine of the 20 products lowered the price per box by more than five percent after the downsizing. For example, Chex Rice 15.6oz. changed to 12.8oz., implying a 21.9% hidden price increase, but the average price per pound for that product only increased by 7.8% after the downsizing. This discrepancy is due to the fact that the average price per box for this product declined by 10.2% in the after-downsizing period.

Finally, Table 2.4 provides average expenditure share information for each of the products. The expenditure share is calculated as the proportion of a good's dollar sales to the dollar sales of all 150 goods in the sample at a city-supermarket-week. Some of the products increased share in the after downsize period (e.g., Cheerios 20oz.), while others saw a decline in expenditure share (e.g., Cheerios 15oz.). For some products, the large swings may be attributable to the decrease

²³There may have been more, however we only focus on a subsample of 65 out of the 150 cereal products in our sample. Out of 65 cereal products, 20 out of 23 General Mills products were downsized.

in price per box. For example, Chex Rice 15.6oz. increased expenditure share by 15.6%, which may be due to the 10.2% drop in price per box that outweighed the hefty downsizing from 15.6 ounces to 12.8 ounces. On the other hand, Cheerios 20oz. witnessed an increase in expenditure share of 22.7% after its downsizing, but its price per box decreased by only 2.8%. We explain later why this phenomenon does not contradict the implications of Proposition 1.

Thus, there appears to be other factors that are influencing the expenditure share of the products in the before to after downsizing periods. This may include not just the good’s own price, but the good’s promotional intensity, the prices of the numerous other cereals in the market, seasonality, and trend. By estimating a full demand system, we take these factors into consideration.

2.5 Empirical Results

In this section, we conduct four empirical investigations to determine the impact of hidden price changes on GM’s cereal business. First, we estimate a basic linear demand model, extending GK’s specification to include prices of all cereals and promotional variables. Second, we estimate demand with the AIDS model to predict expenditure share for the after-downsizing sample period, in order to compare it to the actual expenditure share. We also estimate a difference-in-difference specification for expenditure share to determine how GM’s portfolio of products that were downsized fared as a whole. Finally, we discuss Megerdichian’s (2009) analysis of product downsizing’s effect on GM’s profitability.

2.5.1 Linear Demand

Table 2.5 presents basic linear demand estimates tantamount to GK’s “individual regressions” presented in their Table 4 and discussed above in Section 2. Here, we first focus on GK’s specification as it is a reasonable starting point, given that it’s the only empirical study of hidden price changes to date. For each of the $j = 1, 2, \dots, J = 20$ GM products that underwent downsizing presented in Table

2.3, quantity demanded is described by the data generation process (DGP):

$$q_{jm} = \lambda_j + \alpha_j s_{jm} + \gamma_j p_{jm} + \omega_j T + \lambda_{jt} + \lambda_{jr} + \varepsilon_{jm}, \quad (2.10)$$

where q is quantity purchases of boxes of good j in market m , s is the box size in ounces, p is price measured in dollars per box, T is a trend variable, λ_{jt} represents time fixed effects,²⁴ λ_{jr} represents city-retailer fixed effects, and ε_{jm} captures unobservable variables that drive q_{jm} . A market m is defined as before (see subsection 2.3.1 for details). The OLS estimates presented in Table 2.5 are based on 157 weeks for each of the 121 city-supermarkets.²⁵

As in GK’s analysis, Table 2.5’s results show that good j ’s own price “dominates” the box size variable in explaining quantity demanded (in the sense defined by GK). Of course price and box size are measured in different units, but their t-statistics are unit free, and they show that price has a much “stronger” impact on quantity than box size does; while price is always statistically significant, the box size variable rarely is. Moreover, estimates of α_j are negative except for Lucky Charms 20oz. and Reese’s 14.25oz (both insignificant at the 10% confidence level).

Recall that GK noted that their finding of negative estimates of α_j were “counterintuitive”; yet we maintain that $\hat{\alpha}_j \equiv \partial q_j / \partial s_j < 0$ is in line with consumers not noticing hidden price changes, since when s declines due to downsizing we expect q to increase in a relevant time period because consumers will purchase more boxes since there is less cereal in each box. Moreover, when we concentrate attention on product families, the estimates in Table 2.5 show that this *positive* impact of downsizing on quantity is more pronounced for bigger boxes, as predicted by our model. For example, the coefficients for the Cheerios and the Cheerios Honey Nut families of products are significant at the 10% confidence level only for the bigger size boxes, while they are insignificant for the smaller boxes. In general, for any given family of products in our sample, the p-values for the size variable are always bigger for the smaller size boxes. This observation supports the hypothesis that consumers fail to recognize downsizing of bigger boxes, but not the downsizing of the smaller ones.

²⁴Based on months, not weeks, although the frequency of the data is weekly.

²⁵The number of observations would be $121 \times 157 = 18,997$ but for missing price data.

Table 2.6 presents estimates of the own price effect, γ_{jk} for $j = k$, and the box size effect, α_j , for a more detailed specification of a linear demand DGP given by

$$q_{jm} = \lambda_j + \alpha_j s_{jm} + \sum_k \gamma_{jk} p_{km} + \mathbf{x}_{jm} \theta_j + \mathbf{z}_m \pi_j + \omega_j T + \lambda_{jt} + \lambda_{jr} + \varepsilon_{jm}. \quad (2.11)$$

Prices are given by p_k for $k = 1, 2, \dots, K = 65$ cereal products,²⁶ \mathbf{x} is a vector of good j 's marketing-mix or promotional variables,²⁷ and \mathbf{z} is a vector of income and demographic variables.²⁸ In comparison to GK's specification in Table 2.5, estimates of the own price effect have decreased in magnitude with the full specification, which is due to the omitted variable bias plaguing the simple GK specification. Also, some of the negative estimates of α in Table 2.5 are now positive in Table 2.6, although still statistically insignificant. A positive estimate of the size effect, $\hat{\alpha}_j > 0$, suggests that consumers notice the downsizing and resulting hidden price change and are buying fewer boxes when the size declines. However, note that six of size coefficients are significant at a 10% confidence level, and from those only two are positive. Additionally, for products belonging to a family of products, the significant negative coefficients still associated with the larger boxes.

Consider the estimate of α_j for GM Cheerios 15oz., $\hat{\alpha}_j = -830$. GM Cheerios downsized from 15 ounces to 14 ounces (see Table 2.3), meaning that sales of Cheerios 15/14 oz. are expected to increase by 830 boxes as a result of the downsizing. On the other hand, Cheerios Honey Nut downsized from 27 ounces to 25.25 ounces, resulting in a loss in sales of 250 boxes after the downsizing of that product. Two primary conclusions emerge. First, it appears that the box size variable doesn't explain as much of the variation in quantity demanded as price. This is in line with GK's conclusion that overt price changes are easier to detect

²⁶Including 23 General Mills products (20 of which are the downsized products in Table 2.3), 26 Kellogg products, 9 Post products, 3 Quaker products, and 4 store brand products. See Megerdichian (2010) for a complete list of the 65 products.

²⁷Includes ad only, display only, ad and display, and distribution. See Little (1998) for a detailed description of these variables. We've also included the number of weeks that have elapsed since the last promotion took place within a city-supermarket. See Megerdichian (2009, 2010) and Pesendorfer (2002).

²⁸Includes percentage of population under age 19, percentage of population over age 55, percentage of households with children, percentage of population that is white, income (weekly average wage).

than hidden price changes. Second, some products increased sales (in boxes) after the downsizing, while others did not, as expected.

This approach, although an informative starting point, suffers a few drawbacks. First, even our “better-specified” version of GK’s simple demand analysis is problematic because it is an *ad hoc* linear demand equation that may not be a good approximation for the true DGP for quantity sales. Moreover, properly making causal inferences about the own price effects, γ_{jk} for $j = k$, and the box size effect, α_j , relies on the assumption that we have identified those causal effects, which may be unreasonable due to misspecification or endogeneity. We refer the reader to Hausman (1997), Nevo (2001), and Megerdichian (2009, 2010) for an extensive review of addressing endogeneity in demand estimation. In the next section, we attempt to circumvent these issues with a predictive demand model of expenditure shares to determine whether downsizing was a successful strategy for General Mills.

2.5.2 AIDS Predictions

We use the AIDS specification described previously to predict the expenditure share for the 20 GM goods in question after the downsizing took place.²⁹ Our approach uses the pre-downsizing period to estimate how consumers react to overt price changes, and then compares the forecasted expenditure shares with the actual ones (which take effect after the downsizing). If consumers are fully rational, then the ratio of actual to predicted expenditures shares shouldn’t be statistically different from one, implying that consumers cannot be misled to by implicit price changes.

More specifically, in the first stage of the analysis the AIDS equation for the expenditure share of good j is estimated using only the before-downsizing sample period. In the second stage, the predicted expenditure share, \hat{w}_j , is obtained by evaluating the estimated AIDS equations using the after-downsizing sample period. If a sufficient number of consumers didn’t notice the product downsizing

²⁹For further details on the estimation of the AIDS model, see Deaton and Muellbauer (1980), Chalfant (1987), Green and Alston (1990), Hausman (1997), Hausman and Leonard (2002), and Megerdichian (2010).

and resulting hidden price change for good j , then we expect the actual expenditure share (w_j) to be higher than the predicted AIDS expenditure share (\hat{w}_j) during the after-downsizing sample period. Because we utilize this demand model for predictive purposes only, we do not pursue the estimation of causal effects.

Table 2.7 presents the mean actual expenditure share and the mean predicted expenditure share for each of the 20 GM products for the after-downsizing sample period. Table 2.7 also presents the ratio of actual to predicted expenditure share, the difference, and the corresponding p-value for the difference from a two-sided t-test of no difference between actual and predicted expenditure share. A ratio greater than one means that the actual expenditure share in the after-downsizing sample period is greater than the predicted expenditure share, indicating that the downsizing was successful for that GM product in terms of gaining revenue relative to the other cereal products in the demand system. Conversely, a ratio less than one means that the actual expenditure share of the good is less than the predicted expenditure share, indicating that the downsizing was not successful for that GM product. A few regression diagnostics from the AIDS estimation are also provided.

Table 2.8 presents the same information as Table 2.7, but sorts in descending order based on the expenditure share ratio. The first eleven products, GM Chex Rice 15.6oz. through GM Cookie Crisp 12.25oz., all have ratios greater than one that are statistically significant. It appears that for these eleven products, GM's campaign to downsize during 2007 was successful from a revenue gain perspective. The next four of the GM products, Reese's 14.25oz. through Cheerios 15oz., have actual-predicted mean shares that are either not statistically different from zero or have ratios close enough to one to call into question whether they gained or lost share relative to what is predicted by the AIDS model. Finally, the last five goods, Golden Grahams 13oz. through Cheerios Honey Nut 14oz., have actual-to-predicted ratios that are far enough below one to conclude that these products did not benefit from the downsizing and resulting hidden price change. That is, enough consumers noticed the hidden price change on these five goods, causing them to substitute to other products, so that GM attained a lower expenditure

share relative to what the AIDS equations predicts for the after-downsizing period.

An interesting pattern emerges from the results in Table 2.8 regarding the three groupings of GM products based on the ratio of actual-to-predicted expenditure share. The first eleven products with ratios greater than one (PANEL A) have an average before-downsizing package size of 18.28 ounces (median 20oz.); the middle four products (PANEL B) have an average size of 15.13 ounces (median 14.63oz.); and the last five products (PANEL C) with ratios sufficiently less than one have an average size of 12.6 ounces (median 13oz.). This suggests that the larger-sized products fared better than smaller-sized products from the downsizing, a result consistent with the implications of our extended AIDS model.

To make this point concrete, Table 2.9 presents 10 of the 20 GM products in our sample that have multiple sizes within the same brand-line, including Cheerios (10oz, 15oz., 20oz.), Cheerios Honey Nut (14oz., 20oz., 27oz.), Cinnamon Toast Crunch (14oz., 20.25oz.), and Lucky Charms (14oz., 20oz.). Within each of the four brand groupings, the products are sorted in descending order based on the actual-to-predicted expenditure share ratio. What is of interest here is that the ratio within each product grouping is nearly monotonic in box size, with only a slight exception for 20oz. and 27oz. Cheerios Honey Nut. For example, consider the first grouping of the three Cheerios products. As we go from 20oz. to 15oz. to 10oz. down the list, the expenditure share ratio declines from 1.22 to 0.94 to 0.59, respectively.

This pattern generally holds for each of the four brand groupings, as the larger boxes within the brand enjoy higher ratios, indicating they tended to succeed in gaining expenditure share relative to the smaller boxes of the same brand based on the AIDS predictions. This evidence suggests that the downsizing caused consumers to shift from the smaller boxes of a brand to the larger boxes of the same brand, while consumers of larger boxes missed the implicit price change.

2.5.3 A Diff-in-Diff Approach

The AIDS analysis presented above resulted in 11 products gaining share relative to what the AIDS model predicted, while 9 of the products were either

even or lost share. Here, we take a different approach to determine how GM’s 20 downsized products performed after the downsizing.

We consider a difference-in-difference (diff-in-diff) strategy, which can be particularly useful for natural experiments. The downsizing that took place by GM in July 2007 serves as a “natural” experiment because the impetus for the downsizings and resulting hidden price increases was the dramatic increase in input prices faced by ready-to-eat cereal producers (see Table 2.1 and Figure 2.1). That GM’s marginal costs exogenously increased, causing them to raise price by downsizing, is good reason to believe that the downsizing event serves as a treatment in a natural experiment. As a concomitant, we are hopeful that this approach partly allays Angrist and Pischke’s (2010) concerns regarding causal inference in empirical IO studies.

We consider the following diff-in-diff specification:³⁰

$$w_{krt} = \lambda + \beta_1 G_k + \beta_2 D_t + \beta_3 GD_{kt} + \gamma p_{krt} + \mathbf{x}_{krt}\theta + \mathbf{z}_{rt}\pi + \omega T + \lambda_k + \lambda_r + \lambda_t + \varepsilon_{krt}, \quad (2.12)$$

where k indexes the 65 goods in the sample, r indexes city-retailers, and t indexes week. G_k is an indicator variable that takes on a value of one if good k is one of the 20 GM products that downsized, D_t is an indicator variable that takes on a value of one if week t is during the after-downsize period, and GD_{kt} is the product of G_k and D_t . Finally, the variables p , \mathbf{x} , \mathbf{z} , and T are defined as in the earlier AIDS model, λ_k denotes product fixed effects, λ_r denotes city-retailer fixed effects, and λ_t denotes month fixed effects. Here, w_{krt} is the expenditure share of good k calculated as the proportion of dollar sales of good k in r during t to the dollar sales of all 150 products in the full sample (not just the $K = 65$ products in the AIDS analysis) in r during t .³¹

The diff-in-diff quantity of interest is $\Delta \equiv (w_g^a - w_g^b) - (w_{\tilde{g}}^a - w_{\tilde{g}}^b)$, where g denotes the 20 GM products that underwent downsizing (the “treatment” or “intervention”), \tilde{g} denotes the 45 other products in the sample that are not one

³⁰For an excellent example of this type of identification strategy in IO, see Shepard (1991).

³¹Here, expenditure share is multiplied by 100 for scale purposes in reporting the regression results. E.g., if good k has 1.8 percent share, this is $w_{krt} = 1.8$ in the data, not $w_{krt} = 0.018$.

of the 20 downsized GM products (the “control”), a denotes the after-downsize sample period, and b denotes the before-downsize sample period. The coefficient of interest is β_3 , as this captures the diff-in-diff value of Δ :

$$\begin{aligned}
\Delta &\equiv [w_g^a - w_g^b] - [w_g^a - w_g^b] \\
&= [E(w|G = 1, D = 1, GD = 1) - E(w|G = 1, D = 0, GD = 0)] \\
&\quad - [E(w|G = 0, D = 1, GD = 0) - E(w|G = 0, D = 0, GD = 0)] \\
&= [(\beta_1 + \beta_2 + \beta_3 + \dots) - (\beta_1 + \dots)] - [(\beta_2 + \dots) - (0 + \dots)] \\
&= \beta_3.
\end{aligned}$$

Thus, the estimate of β_3 captures the aggregate net effect of the downsizing for the 20 GM products under investigation during the after-downsizing period, conditional on the additional covariates described above.

Column (6) of Table 2.11 reports the OLS results. For comparison purposes, columns (1) through (3) in Table 2.10 and columns (4) and (5) in Table 2.11 exclude certain variables. The estimate of β_3 is sensitive to specification, as it varies from -0.546 to 0.025 depending on the included covariates and fixed effects. The r -squared and AIC improve substantially with the full specification in column (6). The estimate of β_3 changes from negative to positive when product fixed effects are included in column (6). The estimate of β_3 from column (6) is 0.025 , which means that GM gained about 0.025 percentage points (e.g., from 23% to 23.025%) on its 20 downsized products during the after-downsizing period. In other words, the effect of the downsizing that remains after controlling for confounding is an anemic increase at best, suggesting that, overall, not enough consumers fell for the hidden price change to increase GM’s total expenditure share by a meaningful amount.

Therefore, this analysis, in conjunction with the individual product-by-product AIDS prediction analysis conducted earlier, suggests that some products performed well after the downsizing, while others did not (due perhaps to cannibalization), which ultimately leads to a negligible net aggregate effect for expenditure share. In the next section, we examine the impact of the downsizing on GM’s profitability.

2.5.4 Logit Demand and Profitability

Finally, here we discuss Megerdichian's (2009) analysis of GM's change in profitability due to the downsizing of their 20 products in our sample. He first estimates a logit model of demand to obtain a matrix of demand elasticity estimates before and after the hidden price change (approximately July 2007).³² Consistent with Nevo (2001), he assumes a multiproduct differentiated Nash-Bertrand oligopoly model for the cereal industry, and then proceeds to estimate implied price-cost margins (PCM) for each good j , $(p_j - c_j)/p_j$, for the before-downsizing and after-downsizing sample periods based on the demand estimates. Table 2.12 presents the PCM results for the before and after downsizing periods found in Table XI of Megerdichian (2009). It shows that General Mills increased its profitability by 3.6%; Kellogg declined by 2.8%; Post increased by 4.2%; Quaker increased by 1.9%; and for store brands that were no impacts on profitability.

We note that the estimates of GM's increase in profitability after the downsizing is roughly consistent with external company reports. GM reported that its net sales for "Big G" cereals in their U.S. retail segment increased by 2.7% for the quarter ending November 25, 2007, and that their operating profits for their entire U.S. retail segment grew 1.4% for the 6-month period ending November 25, 2007.^{33,34} The change in the profitability of GM found in Table 2.12 (+3.6%) is roughly in line with the accounting information. The discrepancy is not unexpected for three reasons: (i) The accounting information reported is for all retail channels, not just supermarkets; (ii) the time periods are not consistent; and (iii) the profitability estimates for GM in Table 2.12 is for only 23 of GM's cereal products, not their entire portfolio of products as in the accounting reports. Nevertheless, the 3.6% estimated increase for GM supports the idea that after raising price by downsizing, GM increased their profits to get back to profit maximization, which had been eroded due to the rising input costs in 2006 and early 2007.

³²Extensive attention is paid to identification; in particular, he evaluates standard instrumental variables approaches as well as applying conditional independence assumptions and related instruments.

³³Includes all their U.S. product segments, such as yogurt, snacks, baking products, cereals, and more.

³⁴Source: General Mills SEC Form 10-Q for quarter ended November 25, 2007.

Another result from Table 2.12 that is in line with what has transpired in the cereal industry is that Kellogg's profitability decreased after GM downsized, while all the other firms either increased profitability or did not change. Beginning in 2008, which is beyond our sample, Kellogg also increased the prices of its cereals through downsizing, which none of the other cereal manufacturers have done to date. This action by Kellogg, and the corresponding non-action by the other firms, comports with the margin calculations in Table 2.12. Kellogg's profitability declined during that period because they too faced rising input costs in 2007, yet they didn't commit to a hidden price increase until 2008. General Mills moved first in 2007 with the downsizing, and then Kellogg—the only firm with an estimated drop in profitability—is the only firm to have followed suit and downsized in 2008. Moreover, the margins for the other three manufacturers have either increased or not changed, potentially explaining why they have not followed General Mills and Kellogg with hidden price increases.

2.6 Conclusions

We investigated the hidden price change phenomenon that results from product downsizing in the ready-to-eat cereal industry. General Mills, the second largest manufacturer of cereal in the U.S., increased the price of most of its cereals by decreasing the content of its boxes in July 2007. This change in price was driven by an increase in several key commodity prices that are inputs for producing cereal.

We examined the few studies on hidden price changes and product downsizing. Gourville and Koehler (2004), while limited and flawed in many ways, is the most comprehensive quantitative study on product downsizing to date. Yet it does not adequately address how successful downsizing is from a revenue or profit standpoint. We contribute to the literature by undertaking several empirical studies on a rich set of scanner data, and by presenting a oligopoly model to demonstrate how successful downsizing depends on the probability that consumers will fail to recognize implicit price changes.

From our empirical investigations we conclude that out of the 20 General

Mills products in our sample that downsized, about half the products increased expenditure share in the category by more than what the demand model predicts, indicating a sufficient number of consumers didn't notice the product downsizing for these goods. An alternative, but not excluding interpretation, is that these products cannibalized sales from the other downsized products that became too small for consumers' tastes. In this case, firms should use downsizing in an attempt to foster cannibalization "inside" product families, believing that by restraining themselves from overt price changes their sales wouldn't be cannibalized by other producers's goods. This explains why GM downsized most of its products, and not only the larger boxes.

A key finding is that the products that did benefit from the downsizing tended to be large-sized boxes (by weight), while those that did not benefit from the downsizing tended to be small-sized boxes. One explanation is that consumers are more likely to notice when already small boxes get smaller, and substitute to the larger boxes. Moreover, when examining different size products within the same brand-line (e.g., Cheerios 10oz., 15oz., and 20oz.), the larger-sized box is the one that benefitted the most from the downsizing relative to its smaller-sized sibling. In fact, this pattern held nearly perfectly within each of the four brand groupings. A potential explanation here is that consumers of larger boxes don't notice hidden price changes, and at the same time some consumers of smaller sizes prefer boxes not to be too big nor too small, thus they substitute away from small boxes that have been downsized to large boxes that have been downsized.

Because the downsizing by General Mills in 2007 was due to an exogenous increase in its input costs, it is arguably a natural experiment that lends itself to diff-in-diff estimation. After controlling for a number of confounding factors, using the diff-in-diff approach we find that the downsizing had a negligible effect on the expenditure share for the 20 General Mills products as a whole. While this result may appear to suggest that consumers noticed the hidden price change, given our previous result it is more consistent with the notion that some of General Mills products benefitted because enough consumer didn't notice, while other products did not benefit because consumers did notice, thus leading to a negligible net effect.

Finally, we examined Megerdichian's (2009) investigation of General Mills' change in profitability due to the downsizing. Two of his main findings are consistent with accounting profits and what has transpired in the cereal industry since the end of 2007. His profitability estimates for General Mills before and after the downsizing are approximately consistent with accounting profits reported by the firm. His finding that General Mills' profits increased by nearly 4% after the downsizing is expected if increasing prices was the appropriate profit maximizing action to undertake when they incurred higher input costs in 2006 and 2007. Lastly, the only firm whose profitability declined after General Mills' downsizing was Kellogg, which is the only firm that also committed downsizing starting in 2008, most likely in an attempt to re-attain profit maximization as General Mills did. Unfortunately, our dataset does not extend past 2008 and we cannot show at this time how downsizing by a firm can affect the successfulness of the next player regarding downsizing. In future work we intend to investigate if a previous round of downsizing can cause consumers to be more aware of implicit price changes, making downsizing less effective for the follower than it is for the leader firm.

This chapter is based on *Product Downsizing and Hidden Price Changes in the Ready-to-Eat Cereal Market*, joint work with Aren Megerdichian.

2.7 Tables and Figures




Figure 2.1: Input Commodity Prices (January 2006 = 100)

Table 2.1: Commodity Prices

Year	Petroleum		Corn		Rice		Sugar		Wheat	
2003	46	%	102	%	70	%	92	%	87	%
2004	61	31	109	6	86	23	88	-4	94	7
2005	86	41	96	-12	101	17	90	2	91	-3
2006	103	20	118	24	107	5	95	5	115	26
2007	114	11	159	34	117	10	89	-6	153	33

Source: International Monetary Fund, Energy and Commodities Surveillance Unit.

The first column is a price index that is constructed from monthly data by taking the average for each year. January 2006 is indexed to 100. The second column is the year-over-year percentage change.

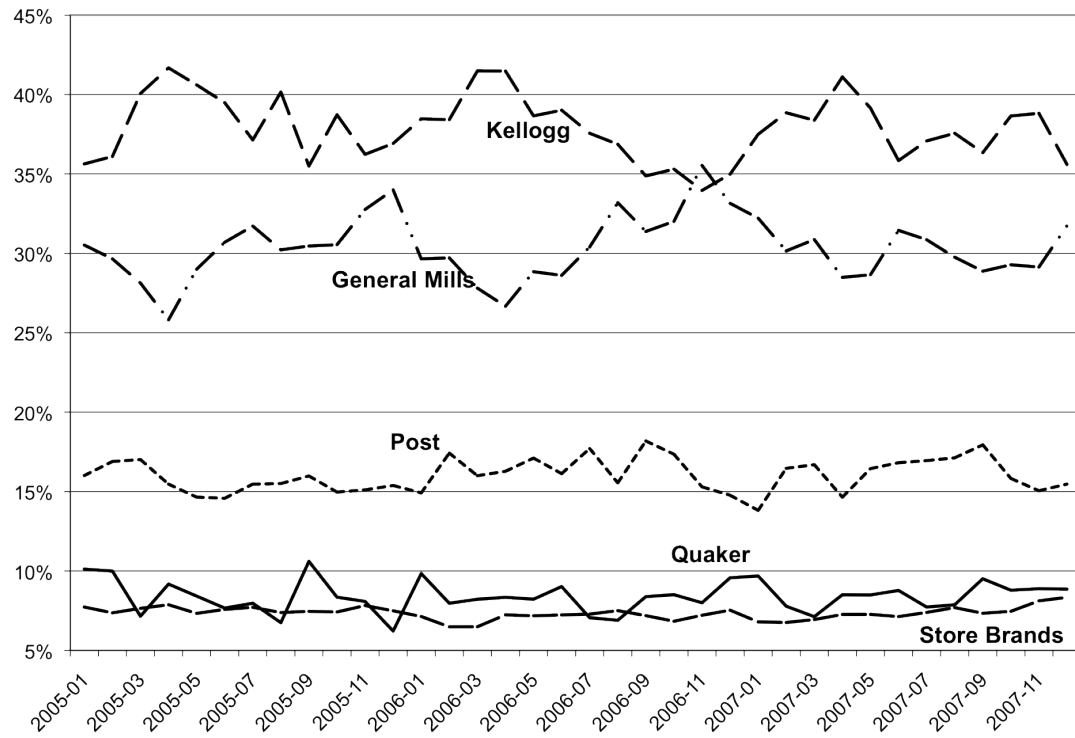


Figure 2.2: Aggregate Market Share




Figure 2.3: Aggregate Price Per Pound

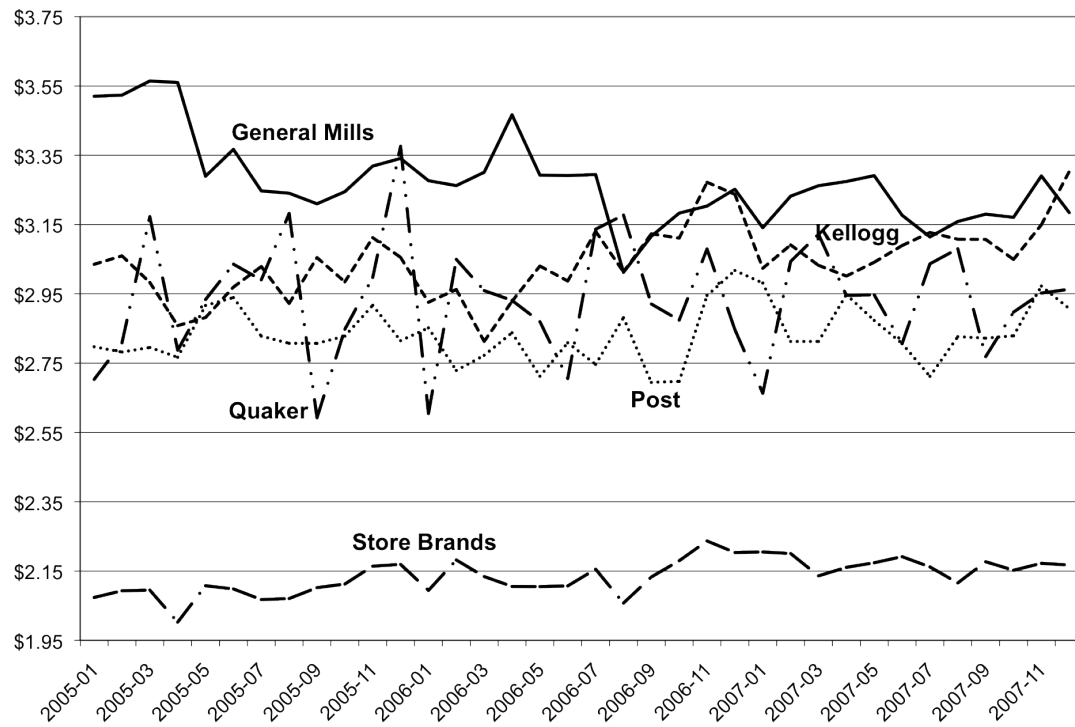


Figure 2.4: Aggregate Price Per Box




Figure 2.5: General Mills Cheerios 15oz. to 14oz.




Figure 2.6: Possible Substitution Patterns Resulting from Downsizing

Table 2.2: Aggregate Market Shares and Average Prices

	% Market Share - DS	% Market Share - PS	Number of SKUs	Wtd. Avg. Price/Lb. (\$)	% Above/Below Total Avg. Price
Kellogg	37.9	37.9	55	2.95	-0.1
General Mills	34.9	30.3	45	3.40	15.2
Post	14.7	16.0	25	2.71	-8.5
Quaker	7.6	8.4	13	2.66	-10.2
Store Brands	5.0	7.3	12	2.01	-31.9
Total	100	100	150	2.96	-

Source: IRI scanner data. Based on 121 city-supermarket chains, 2005-2007. SKU is "stock keeping unit." Market Share - DS is based on dollar sales. Market Share - PS is based on sales in pounds of cereal.

Table 2.3: General Mills Downsizing

GM Products:	Date of downsizing	Ounces per box		Implied price Δ (%)	Price per pound			
		bef	aft		Δ (%)	bef	aft	Δ (%)
Cheerios 10oz	07/19/07	10	8.9	-11.0	12.4	4.73	5.29	11.9
Cheerios 15oz	07/19/07	15	14	-6.7	7.1	3.46	3.80	9.8
Cheerios 20oz	07/11/07	20	18	-10.0	11.1	3.53	3.80	7.6
Cheerios Frosted 20.25oz	07/22/07	20.25	17.2	-15.1	17.7	3.02	3.64	20.5
Cheerios Honey Nut 14oz	07/24/07	14	12.25	-12.5	14.3	3.72	4.25	14.0
Cheerios Honey Nut 20oz	07/13/07	20	17	-15.0	17.6	3.56	4.17	17.1
Cheerios Honey Nut 27oz	07/12/07	27	25.25	-6.5	6.9	3.13	3.40	8.7
Chex Rice 15.6oz	07/20/07	15.6	12.8	-17.9	21.9	3.63	3.91	7.8
Cinnamon Toast Crunch 14oz	08/04/07	14	12.8	-8.6	9.4	3.60	3.73	3.8
Cinnamon Toast Crunch 20.25oz	07/20/07	20.25	17	-16.0	19.1	3.51	3.91	11.4
Cocoa Puffs 13.75oz	07/25/07	13.75	11.8	-14.2	16.5	3.68	3.89	5.4
Cookie Crisp 12.25oz	08/04/07	12.25	11.25	-8.2	8.9	5.01	5.18	3.4
Golden Grahams 13oz	08/13/07	13	12	-7.7	8.3	4.10	3.88	-5.3
Lucky Charms 14oz	07/22/07	14	11.5	-17.9	21.7	3.84	4.52	18.0
Lucky Charms 20oz	07/19/07	20	16	-20.0	25.0	3.80	4.42	16.4
Oatmeal Crisp Raisin 19.25oz	08/08/07	19.25	18	-6.5	6.9	3.35	3.52	5.0
Reeses PNT BTR Puffs 14.25oz	07/22/07	14.25	13	-8.8	9.6	4.08	4.13	1.3
Total Whole Grain 12oz	08/05/07	12	10.6	-11.7	13.2	4.83	5.11	5.8
TRIX 12oz	08/06/07	12	10.7	-10.8	12.1	4.13	4.23	2.5
Wheaties 18oz	07/23/07	18	15.6	-13.3	15.4	3.53	4.11	16.4

Source: IRI scanner data. Based on 121 city-supermarket chains, 2005-2007.

Table 2.4: General Mills Downsizing, Table 2.3 Continued

GM Products:	Price per box			Expenditure share		
	bef	aft	Δ (%)	bef	aft	Δ (%)
Cheerios 10oz	2.95	2.96	0.21	0.02	0.02	0.10
Cheerios 15oz	3.25	3.34	2.74	0.03	0.02	-17.20
Cheerios 20oz	4.41	4.29	-2.78	0.02	0.02	22.72
Cheerios Frosted 20.25oz	3.82	3.97	3.89	0.01	0.01	-1.36
Cheerios Honey Nut 14oz	3.26	3.27	0.40	0.03	0.03	1.43
Cheerios Honey Nut 20oz	4.45	4.47	0.37	0.02	0.02	-1.45
Cheerios Honey Nut 27oz	5.28	5.38	1.83	0.01	0.02	8.90
Chex Rice 15.6oz	3.54	3.18	-10.18	0.01	0.01	15.57
Cinnamon Toast Crunch 14oz	3.20	3.00	-6.05	0.01	0.01	-2.08
Cinnamon Toast Crunch 20.25oz	4.44	4.21	-5.10	0.01	0.01	2.44
Cocoa Puffs 13.75oz	3.21	2.91	-9.47	0.01	0.01	4.04
Cookie Crisp 12.25oz	3.84	3.67	-4.54	0.01	0.01	-23.39
Golden Grahams 13oz	3.34	2.93	-12.15	0.01	0.01	14.49
Lucky Charms 14oz	3.39	3.30	-2.69	0.01	0.01	-0.26
Lucky Charms 20oz	4.75	4.50	-5.33	0.01	0.01	8.61
Oatmeal Crisp Raisin 19.25oz	4.03	3.98	-1.23	0.00	0.00	-10.59
Reeses PNT BTR Puffs 14.25oz	3.64	3.37	-7.40	0.01	0.01	2.77
Total Whole Grain 12oz	3.62	3.43	-5.17	0.01	0.01	-15.73
TRIX 12oz	3.15	2.87	-8.70	0.01	0.01	-7.68
Wheaties 18oz	3.97	4.06	2.11	0.01	0.01	-3.64

Source: IRI scanner data. Based on 121 city-supermarket chains, 2005-2007.

Table 2.5: Single Price Linear Demand - GK's Specification

Dependent var: qty boxes	Price per Box		Box Size		Regression Stats		
	coeff*	tstat	coeff	tstat	obs*	r^2	RMSE*
Cheerios 10oz	-3.66	-27.53	-236.9	-1.14	18.76	0.47	2.655
Cheerios 15oz	-3.40	-40.58	-806.7	-3.45	18.99	0.52	3.162
Cheerios 20oz	-1.61	-22.03	-71.4	-1.83	19.00	0.47	1.480
Cheerios Frosted 20.25oz	-0.57	-27.08	-6.2	-0.80	17.65	0.48	0.536
Cheerios Honey Nut 14oz	-4.19	-40.93	-68.4	-0.46	19.00	0.50	3.845
Cheerios Honey Nut 20oz	-2.04	-21.33	-63.8	-1.92	18.98	0.46	1.828
Cheerios Honey Nut 27oz	-0.94	-17.16	-72.1	-1.74	17.51	0.43	1.151
Chex Rice 15.6oz	-0.50	-34.97	-6.0	-0.71	18.02	0.62	0.451
Cinnamon Toast Crunch 14oz	-2.64	-40.59	-8.1	-0.20	19.00	0.41	2.790
Cinnamon Toast Crunch 20.25oz	-1.01	-28.21	-15.4	-0.64	18.99	0.44	1.088
Cocoa Puffs 13.75oz	-1.49	-30.52	-19.6	-0.53	18.64	0.40	1.583
Cookie Crisp 12.25oz	-1.27	-29.11	-144.0	-3.53	18.89	0.43	1.470
Golden Grahams 13oz	-1.14	-31.86	-20.6	-0.58	17.60	0.39	1.145
Lucky Charms 14oz	-1.98	-43.36	-21.5	-0.49	19.00	0.44	2.271
Lucky Charms 20oz	-0.90	-26.73	11.8	0.91	18.84	0.44	0.952
Oatmeal Crisp Raisin 19.25oz	-0.40	-20.11	-17.0	-0.88	17.75	0.51	0.385
Reeses PNT BTR Puffs 14.25oz	-1.10	-38.61	3.6	0.09	18.87	0.43	1.224
Total Whole Grain 12oz	-0.91	-16.52	-143.3	-2.00	18.46	0.38	1.005
TRIX 12oz	-1.44	-30.79	-6.0	-0.16	18.73	0.37	1.509
Wheaties 18oz	-0.64	-20.02	-18.5	-1.79	18.42	0.53	0.542

T-stats and p-values are based on White (1980) heteroskedasticity-robust standard errors. * Values expressed in thousands.

Table 2.6: Full Specification, Linear Demand

Dependent var: qty boxes	coeff*	tstat	pval	coeff	tstat	pval	obs*	r^2	RMSE*
Cheerios 10oz	-2.83	-17.46	0	0.5	0.00	1.00	8.25	0.59	2.502
Cheerios 15oz	-3.08	-19.24	0	-830.3	-2.69	0.01	8.25	0.61	3.047
Cheerios 20oz	-1.25	-14.33	0	43.6	0.91	0.37	8.25	0.62	1.286
Cheerios Frosted 20.25oz	-0.48	-14.87	0	-32.5	-2.52	0.01	8.25	0.54	0.557
Cheerios Honey Nut 14oz	-3.02	-20.46	0	62.9	0.41	0.68	8.25	0.59	3.380
Cheerios Honey Nut 20oz	-1.25	-13.50	0	39.9	1.26	0.21	8.25	0.63	1.431
Cheerios Honey Nut 27oz	-0.90	-13.98	0	142.7	2.49	0.01	8.25	0.58	1.083
Chex Rice 15.6oz	-0.45	-19.10	0	-7.7	-0.62	0.54	8.25	0.69	0.437
Cinnamon Toast Crunch 14oz	-2.03	-20.63	0	34.3	0.61	0.54	8.25	0.54	2.168
Cinnamon Toast Crunch 20.25oz	-0.52	-14.33	0	34.1	1.50	0.14	8.25	0.63	0.664
Cocoa Puffs 13.75oz	-1.15	-18.47	0	-109.2	-2.78	0.01	8.25	0.52	1.332
Cookie Crisp 12.25oz	-0.96	-16.62	0	-66.9	-1.31	0.19	8.25	0.57	1.157
Golden Grahams 13oz	-0.91	-17.68	0	62.1	2.15	0.03	8.25	0.50	1.091
Lucky Charms 14oz	-1.48	-21.01	0	32.3	0.54	0.59	8.25	0.57	1.702
Lucky Charms 20oz	-0.57	-14.42	0	-15.6	-1.29	0.20	8.25	0.64	0.579
Oatmeal Crisp Raisin 19.25oz	-0.31	-12.52	0	10.5	0.57	0.57	8.25	0.61	0.347
Reeses PNT BTR Puffs 14.25oz	-0.88	-16.56	0	43.9	0.70	0.48	8.25	0.47	1.239
Total Whole Grain 12oz	-0.74	-8.53	0	17.6	0.25	0.80	8.25	0.46	0.847
TRIX 12oz	-1.21	-21.04	0	-21.9	-0.79	0.43	8.25	0.54	1.179
Wheaties 18oz	-0.55	-12.40	0	-42.0	-2.82	0.01	8.25	0.59	0.552

T-stats and p-values are based on White (1980) heteroskedasticity-robust standard errors. * Values expressed in thousands.

Table 2.7: AIDS Prediction of Expenditure Share

Product	Mean Expenditure Share (After)				Regression Stats			
	actual	predicted	ratio	diff	p-val	obs	r^2	RMSE
Cheerios 10oz	2.95	5.02	0.59	-2.07	0.000	7152	0.696	0.011
Cheerios 15oz	3.73	3.98	0.94	-0.26	0.009	7145	0.759	0.013
Cheerios 20oz	3.28	2.69	1.22	0.59	0.000	7114	0.749	0.007
Cheerios Frosted 20.25oz	0.89	0.76	1.17	0.13	0.000	7160	0.708	0.003
Cheerios Honey Nut 14oz	4.15	7.23	0.57	-3.07	0.000	7224	0.777	0.014
Cheerios Honey Nut 20oz	2.61	1.89	1.39	0.73	0.000	7125	0.785	0.008
Cheerios Honey Nut 27oz	2.69	2.06	1.30	0.62	0.000	7100	0.757	0.007
Chex Rice 15.6oz	1.23	0.06	21.79	1.18	0.000	7143	0.769	0.004
Cinnamon Toast Crunch 14oz	2.33	3.87	0.60	-1.53	0.000	7292	0.768	0.010
Cinnamon Toast Crunch 20.25oz	1.61	1.29	1.25	0.32	0.000	7157	0.756	0.005
Cocoa Puffs 13.75oz	1.37	0.82	1.68	0.55	0.000	7178	0.726	0.005
Cookie Crisp 12.25oz	1.21	1.06	1.14	0.15	0.002	7231	0.791	0.005
Golden Grahams 13oz	1.07	1.33	0.81	-0.25	0.000	7245	0.704	0.005
Lucky Charms 14oz	2.09	1.74	1.20	0.35	0.000	7143	0.737	0.008
Lucky Charms 20oz	1.51	0.62	2.43	0.89	0.000	7137	0.728	0.005
Oatmeal Crisp Raisin 19.25oz	0.68	0.73	0.94	-0.05	0.001	7240	0.716	0.002
Reeses PNT BTR Puffs 14.25oz	1.40	1.38	1.01	0.02	0.755	7159	0.619	0.007
Total Whole Grain 12oz	0.84	1.32	0.63	-0.48	0.000	7241	0.682	0.003
TRIX 12oz	1.30	1.38	0.94	-0.08	0.076	7231	0.759	0.005
Wheaties 18oz	1.07	0.78	1.38	0.29	0.000	7165	0.709	0.003

Table 2.8: AIDS Prediction of Expenditure Share, Sorted by Ratio

Product:	Mean Expenditure Share (After)				
	actual (%)	predict (%)	ratio	diff (%)	p-val
Chex Rice 15.6oz	1.23	0.06	21.79	1.18	0.000
Lucky Charms 20oz	1.51	0.62	2.43	0.89	0.000
Cocoa Puffs 13.75oz	1.37	0.82	1.68	0.55	0.000
Cheerios Honey Nut 20oz	2.61	1.89	1.39	0.73	0.000
Wheaties 18oz	1.07	0.78	1.38	0.29	0.000
Cheerios Honey Nut 27oz	2.69	2.06	1.30	0.62	0.000
Cinnamon Toast Crunch 20.25oz	1.61	1.29	1.25	0.32	0.000
Cheerios 20oz	3.28	2.69	1.22	0.59	0.000
Lucky Charms 14oz	2.09	1.74	1.20	0.35	0.000
Cheerios Frosted 20.25oz	0.89	0.76	1.17	0.13	0.000
Cookie Crisp 12.25oz	1.21	1.06	1.14	0.15	0.002
Reeses PNT BTR Puffs 14.25oz	1.40	1.38	1.01	0.02	0.755
TRIX 12oz	1.30	1.38	0.94	-0.08	0.076
Oatmeal Crisp Raisin 19.25oz	0.68	0.73	0.94	-0.05	0.001
Cheerios 15oz	3.73	3.98	0.94	-0.26	0.009
Golden Grahams 13oz	1.07	1.33	0.81	-0.25	0.000
Total Whole Grain 12oz	0.84	1.32	0.63	-0.48	0.000
Cinnamon Toast Crunch 14oz	2.33	3.87	0.60	-1.53	0.000
Cheerios 10oz	2.95	5.02	0.59	-2.07	0.000
Cheerios Honey Nut 14oz	4.15	7.23	0.57	-3.07	0.000

Table 2.9: AIDS Predictions, Multiple Sizes of Same Brand

Product:	Mean Expenditure Share (After)				
	actual	predict	ratio	diff	p-val
Cheerios 20oz	3.28	2.69	1.22	0.59	0.00
Cheerios 15oz	3.73	3.98	0.94	-0.26	0.01
Cheerios 10oz	2.95	5.02	0.59	-2.07	0.00
Cheerios Honey Nut 20oz	2.61	1.89	1.39	0.73	0.00
Cheerios Honey Nut 27oz	2.69	2.06	1.30	0.62	0.00
Cheerios Honey Nut 14oz	4.15	7.23	0.57	-3.07	0.00
Cinnamon Toast Crunch 20.25oz	1.61	1.29	1.25	0.32	0.00
Cinnamon Toast Crunch 14oz	2.33	3.87	0.60	-1.53	0.00
Lucky Charms 20oz	1.51	0.62	2.43	0.89	0.00
Lucky Charms 14oz	2.09	1.74	1.20	0.35	0.00

Table 2.10: Diff-in-Diff Expenditure Share

Dep. Var.:	(1)		(2)		(3)	
	coeff	p-val	coeff	p-val	coeff	p-val
Expenditure Share x 100						
1[Downsize]	0.547	0.000	0.537	0.000	0.378	0.000
1[GM 20 Products]	0.338	0.000	0.468	0.000	0.355	0.000
1[Downsize] x 1[GM 20 Products]	-0.546	0.000	-0.447	0.000	-0.351	0.000
Price			-0.272	0.000	-0.115	0.000
Ad Only					0.162	0.000
Display Only					1.706	0.000
Ad and Display					3.132	0.000
Distribution					1.798	0.000
Trend						
Month Fixed Effects		n		n		n
City-Retailer Fixed Effects		n		n		n
Product Fixed Effects		n		n		n
Other Covariates		n		n		n
N	1,203,794		1,203,794		1,203,794	
R-Sqd	0.031		0.121		0.462	
RMSE	0.909		0.866		0.700	
AIC	3,187,593		3,070,339		2,558,204	
AIC/N	2.648		2.551		2.125	

P-values are based on White (1980) heteroskedasticity-robust standard errors. "Other Covariates" includes weeks elapsed since last promotion, percentage of population under age 19, percentage of population over age 55, percentage of households with children, percentage of population that is white, income (weekly wage).

Table 2.11: Diff-in-Diff Expenditure Share, Table 2.10 Continued

Dep. Var.:	(4)		(5)		(6)	
	coeff	p-val	coeff	p-val	coeff	p-val
Expenditure Share x 100						
1[Downsize]	0.408	0.00	0.394	0.00	0.083	0.00
1[GM 20 Products]	0.349	0.00	0.350	0.00	1.399	0.00
1[Downsize] x 1[GM 20 Products]	-0.344	0.00	-0.326	0.00	0.025	0.00
Price	-0.115	0.00	-0.123	0.00	-0.306	0.00
Ad Only	0.165	0.00	0.191	0.00	0.114	0.00
Display Only	1.710	0.00	1.655	0.00	1.390	0.00
Ad and Display	3.136	0.00	3.122	0.00	2.692	0.00
Distribution	1.798	0.00	2.272	0.00	1.411	0.00
Trend	-0.001	0.00	-0.001	0.00	-0.001	0.00
Month Fixed Effects		y		y		y
City-Retailer Fixed Effects		n		y		y
Product Fixed Effects		n		n		y
Other Covariates		n		n		y
N	1,203,794		1,203,794		1,203,794	
R-Sqd	0.427		0.445		0.595	
RMSE	0.699		0.689		0.588	
AIC	2,55,122		2,517,900		2,137,608	
AIC/N	2.123		2.092		1.776	

P-values are based on White (1980) heteroskedasticity-robust standard errors. "Other Covariates" includes weeks elapsed since last promotion, percentage of population under age 19, percentage of population over age 55, percentage of households with children, percentage of population that is white, income (weekly wage).

Table 2.12: Price-Cost Margins

Brand:	Before (%)	After (%)	Change (%)
General Mills	66.61	70.19	3.58
Kellogg	73.01	70.18	-2.82
Post	60.21	64.40	4.20
Quaker	60.85	62.72	1.88
Store Brand	56.75	56.76	0.01
All 64 Products	67.32	68.18	0.86

Source: Megerdichian (2009), Table XI.

Chapter 3

Economic Regulation and Asymmetry in Ex-Post Stock Returns

3.1 Introduction

The regulation of an economic activity or industry is generally motivated on efficiency grounds, i.e., as a response from a benevolent central planner to a market failure, such as a natural monopoly. In those situations, a formally constituted regulatory institution, acting on behalf of the public interest, tries to emulate the conditions of a long-run competitive market equilibrium, in which industry participants earn a rate of return on their investment compatible with their cost of capital. There are certainly other reasons, not necessarily welfare enhancing or based on the public interest, that also gives rise to regulation. Stigler (1971) and Peltzman (1976) for instance suggested that regulatory institutions are the outcome of special interest groups and lobbies competing for political influence, or, more concretely, for the allocation and distribution of rents. The capture theory of regulation is in sharp contrast with the so-called public interest theory also in trying to answer why regulation ever appears in the first place.

In either the benevolent or the rent-seeking case, the actions and policies of a regulator have a direct and profound impact on the profitability and risk of regulated entities. Peltzman argued that the regulatory process attenuates profit fluctuations, so the risk of the firm decreases when regulation becomes more stringent. This “buffering” hypothesis has been interpreted as a restriction on the regulated firm’s CAPM beta and has been tested by several authors, e.g. Norton (1985) and Davidson, Rangan and Rosenstein (1997), becoming a generally accepted view at least for low-powered forms of regulation. Incentive-based, high-powered regulatory systems such as the price-cap seem to impose more risk on the firm, as claimed for instance by Alexander, Mayer and Weeds (1996) and Grout and Zalewska (2006) who find a positive relation between contract power (incentives for cost reduction) and beta.

One particular aspect that has not received much attention in the literature is the effect of regulation on the third moment (asymmetry) of returns. In some cases the authority may be involved in systematically generating significant (positive or negative) returns for the regulated firm, and these rents would eventually be priced into the company’s stock. There are at least three different regulatory

phenomena that can potentially affect the symmetry of the distribution of profits of regulated firms. First, there is the possibility that the regulator is captured by some interest group aligned with the industry under supervision, or by the government currently in office. The influence of these groups in the workings of the regulator is likely to generate positive returns for the regulated entities if the pressure group is linked to the industry, or negative returns otherwise. Examples of this sort of situation abound in the literature, the classic reference being cab regulation in New York City (Viscusi, Harrington and Vernon, 2005).

Second, regulated firms may be subject to expropriatory actions from the government, which vary from intentional delays in the rate revision process to a full-blown nationalization of assets without fair or timely compensation. These events will, most of the time, be accompanied by some sort of justification from the authorities trying to characterize them as legitimate actions taken in defense of the public interest. Regardless whether these events are indeed what the authorities claim them to be, the fact is that changes in investors expectations regarding the probability of such events are enough to cause some selling pressure on the firm's stock. This effect can become recurrent if the government fails to change those expectations, eventually altering the shape of the *ex-post* stock return distribution.

Third, in industries under the used-and-useful principle of valuation of the rate base, there is also a permanent downside risk for the firm every time the operating assets are scrutinized by the regulator. In the scheduled hearings for rate revision, some assets may be disallowed from the rate base when their acquisition is considered imprudent by the regulator, and no compensation will be due from such a discretionary action. While this practice is important to mitigate the well known Averch-Johnson effect of over-investment in rate-of-return regulated industries, their impact on the risk and in particular the symmetry of returns is clear. See for instance Kolbe, Tye and Myers (1993) for a richer treatment of this issue in the U.S. natural gas sector.

In any case, unless one can clearly identify which of the above mentioned phenomena is predominant in a given industry, which is certainly a very difficult task, it is generally not possible to determine, *ex-ante*, the combined effect of

those actions on the asymmetry of returns of the regulated entities. But they all represent additional factors that affect the symmetry of returns and are shared only, in a lesser or greater extent, by regulated firms. In this paper we wish to test the hypothesis that regulated stocks exhibit returns that are “more asymmetric” than those of a control group formed by non-regulated stocks. To carry such a task we have to deal with the problem of (i) correctly measuring asymmetry in stock returns, and (ii) claiming that the cross-sectional distribution of return asymmetry differ for regulated and non-regulated stocks.

In order to formally state our hypothesis, consider the population of all traded stocks in a given market, and let $S \in \mathbb{R}$ denote some measure of the asymmetry in the distribution of realized returns. Let F_S denote the distribution of this asymmetry measure in the cross-section of stocks, such that $F_S(s)$ is the mass of companies with $S \leq s$, i.e. $N^{-1} \sum_{i=1}^N \mathbf{1}\{s_i \leq s\}$ where $\mathbf{1}\{\cdot\}$ is the indicator function. In the most general case, one could write the joint distribution $F_{S,\mathbf{X}}(s, x_1, \dots, x_k)$ where $X_j, j = 1, \dots, k$ are variables that might be related with S in the population. We concentrate on a simple case with a single explanatory variable. We write $F_{S,R}(s, r)$, where R is a categorical variable indicating whether the firm is regulated ($r = 1$) or not ($r = 0$). Finally, let $f_{S,R}(s, r)$, $f_S(s)$ and $f_R(r)$ denote the joint and marginal density functions associated with $F_{S,R}(s, r)$. In section 3.3.1 we investigate whether the distribution of asymmetry across firms is different among regulated and unregulated ones by testing the null

$$H_0 : f_{S,R}(s, r) = f_S(s)f_R(r), \quad (3.1)$$

for almost all $s \in \mathbb{R}$ and all $r \in \{0, 1\}$. As we do not want to impose any *a priori* structure on $f_{S,R}$, we use a number of non-parametric approaches to test (3.1) against the general alternative $H_A : f_{S,R}(s, r) \neq f_S(s)f_R(r)$ for some $s \in \mathbb{R}$ or $r \in \{0, 1\}$. The first is a test based on 2×2 contingency tables formed after an appropriate dichotomization of S . One advantage of this approach is that we can provide a corrected p -value that explicitly takes into account the fact that we are using estimates of asymmetry in lieu of actual observed data. The exact p -value under this situation will depend on the size and power of each asymmetry test, and the applied correction yields valid inference under measurement errors in S .

We also perform a test of (3.1) based on estimates of the conditional distribution $F_{S|R}(s|r)$, testing whether $H'_0 : F_{S|R}(s|r) = F_S(s)$ for almost all $s \in \mathbb{R}$ and all $r \in \{0, 1\}$ against a general alternative $H'_A : F_{S|R}(s|r) \neq F_S(s)$ for some $s \in \mathbb{R}$ or $r \in \{0, 1\}$. As a matter of fact it is easier to test H'_0 indirectly through the necessary condition

$$F_{S|R}(s|r=0) = F_{S|R}(s|r=1), \quad (3.2)$$

and to that end we use three different approaches. First we apply the Wald-Wolfowitz runs test of randomness on the sequence of zeros and ones in $\{r_i\}$ ordered by the corresponding asymmetry measure \widehat{s}_i . Second, we use the Kolmogorov-Smirnov, Cramer-von Mises and Anderson-Darling tests for equality of distributions, which are based on the empirical CDF (i.e. non-smooth) estimates of the functions in (3.2). We use the improved versions of these tests, developed by Zhang (2006), which have greater power against any differences in the distributions, including location, scale and shape. Lastly, we use kernel smoothing methods to estimate the distributions in (3.2) and a distance measure among them to test whether they differ.

The hypothesized effect of regulation on return asymmetry can imply other statistical features on the conditional distribution $F_{S|R}(s|r)$ that can be tested. If regulation increases the likelihood of asymmetric returns in either direction, then $F_{S|R}(s|r=1)$ is expected to show a larger spread and we may test (3.2) against a scale alternative of the form

$$H'_S : F_{S|R}(s|r=0) = F_{S|R}(\nu s|r=1) \forall s \in \mathbb{R} \text{ and some } \nu \geq 1, \quad (3.3)$$

a one-sided test that can be implemented with the Ansari-Bradley or the Mood statistics, for instance. On the other hand, if we consider any strictly positive measure \mathcal{S} that does not discriminate among right or left asymmetry, then the density of \mathcal{S} conditional on $R=1$ is expected to be a positive location shift,

$$H'_L : F_{S|R}(s|r=0) = F_{S|R}(s-\theta|r=1) \forall s \in \mathbb{R}^+ \text{ and some } \theta \geq 0, \quad (3.4)$$

again a one-sided test. We use the Wilcoxon Rank-Sum statistic to test (3.2) against this specific alternative.

While we can establish a connection between our conjectured effect of regulation on return asymmetry, and formulate specific predictions that can be falsified, there are two major empirical challenges in testing the above hypotheses. First, asymmetry must be estimated, so our fundamental piece of information is measured with error. In order to alleviate this problem, we rely on several nonparametric methods which are robust to the sampling error embedded in the asymmetry estimates.

Second, the extent to which we can argue whether regulation affects asymmetry is conditional to the sample used and no general causal relation can be established. We simply cannot have a group of comparable companies, operating in the same lines of businesses and with the same characteristics as the regulated ones but lacking only the regulatory restrictions, so that the confounding factors are controlled for.

3.2 Measuring Asymmetry in Financial Returns

First we need to establish some notation and discuss the model of returns that will be used throughout the paper. Let y_t denote the (log) return of a given stock in period t and \mathcal{F}_t the σ -algebra generated by the information available at time t . The traditional model for y_t in empirical applications is

$$y_t = \mu(\theta, \mathcal{F}_t) + \eta_t, \quad (3.5a)$$

$$\eta_t = h(\lambda, \mathcal{F}_t)\epsilon_t, \quad (3.5b)$$

where ϵ_t is an i.i.d. innovation with zero mean and unit variance, $\mu_t \equiv \mu(\theta, \mathcal{F}_t)$ and $h_t \equiv h(\lambda, \mathcal{F}_t)$ are the mean and standard deviation of y_t conditional on \mathcal{F}_t , θ and λ are vectors of parameters and $\eta_t \equiv y_t - \mu_t = h_t\epsilon_t$ is the (possibly heteroscedastic) residual. In Appendix B we show the implications of model (3.5) for the moments of y_t , and how the conditional asymmetry of y_t is implied by the unconditional asymmetry of the innovations ϵ_t .

The concept of symmetry of a random variable can be formalized as follows:

Definition A real random variable Y with c.d.f. F_Y and support in \mathbb{R} is said to

have a center of symmetry in $\phi \in \mathbb{R}$ if $F_Y(\phi + a) = 1 - F_Y(\phi - a)$ for all $a \in \mathbb{R}$.

An immediate consequence of this definition is that $F_Y(\phi) = 1/2$, so ϕ must also be the median of Y . The converse is obviously not necessarily true, as the median does not need to be the center of symmetry, which may not exist for a given random variable. In the context of financial data we are usually interested whether the return process $\{Y_t\}$ is symmetric with regards to its conditional (or unconditional) mean, and in this case the definition above is equivalent as saying that $Z_t \equiv Y_t - \mathbb{E}[Y_t|\mathcal{F}_t]$ has the same distribution as $-Z_t$, or $F_{Z_t}(z) = 1 - F_{Z_t}(-z)$ and, if the p.d.f. exists, $f_{Z_t}(z) = f_{Z_t}(-z)$ for all $z \in \mathcal{Y}$. This definition is fairly general and does not yield a single or obvious index of the degree of asymmetry that Y_t might exhibit. As a result, several measures of asymmetry have been proposed. The aim of this section is to review some recent developments that specifically address the statistical features of financial returns. In particular, if a process $\{Y_t\}$ has finite third moment, the most popular measure of asymmetry is the *skewness coefficient*

$$S(Y_t) = \frac{\mathbb{E}[(Y_t - \mathbb{E}Y_t)^3]}{\mathbb{E}[(Y_t - \mathbb{E}Y_t)^2]^{3/2}}. \quad (3.6)$$

In a model such as (3.5), we show in Appendix B how the conditional mean μ_t , the conditional variance h_t and the innovation process ϵ_t contribute to the unconditional skewness of Y_t :

$$S(Y_t) = \frac{\mathbb{E}[(\mu_t - \mathbb{E}\mu_t)^3] + 3 \text{Cov}(\mu_t, \eta_t^2) + \mathbb{E}h_t^3 \mathbb{E}\epsilon_t^3}{(\text{Var} \mu_t + \mathbb{E}h_t^2)^{3/2}}. \quad (3.7)$$

In (3.7) we see that even when the conditional mean has a negligible effect (as in the case of most stocks), the unconditional skewness of Y_t will depend on the skewness of the innovation process and on the conditional variance h_t . However, if in (3.6) we consider expectations conditioned on the information available at the previous period (\mathcal{F}_{t-1}), the conditional skewness of Y_t ,

$$S_{t-1}(Y_t) = \frac{\mathbb{E}_{t-1}[(Y_t - \mathbb{E}_{t-1}[Y_t])^3]}{\mathbb{E}_{t-1}[(Y_t - \mathbb{E}_{t-1}[Y_t])^2]^{3/2}} = \frac{h_t^3 \mathbb{E}_{t-1}\epsilon_t^3}{(h_t^2 \mathbb{E}_{t-1}\epsilon_t^2)^{3/2}} = \mathbb{E}\epsilon_t^3, \quad (3.8)$$

is related only to the unconditional third moment of the innovation process. In this paper we implement methods for measuring both the unconditional and conditional

asymmetry of returns, and (3.7) and (3.8) above shows that we need to estimate a fully specified (and yet flexible) ARMA-GARCH model for the returns of each stock in our sample in order to understand where possible differences in both concepts of asymmetry might be coming from. We give further details on our modeling strategy for returns at the end of this section.

Cramr (1946) shows that, for i.i.d. normal data, the estimator \widehat{S} of (3.6) based on sample moments and under the null of symmetry is asymptotically normal. More specifically,

$$\sqrt{T} \widehat{S} \xrightarrow{d} \mathcal{N}(0, 6), \quad (3.9)$$

where T is the sample size. Despite its widespread use, \widehat{S} has some serious drawbacks when applied to financial returns that are not always considered in empirical applications. First, it is a moment-based measure and some financial assets are so heavy-tailed that their third moments might not exist.

Second, the statistic is centered on the unconditional mean of the series, which may also not be the best choice of location parameter for all assets. Stocks, for example, tend to exhibit some variation in their conditional mean, and a perfectly (conditionally) symmetric return series may be mistakenly interpreted as a skewed one by (3.6).

Third, the asymptotic variance in (3.9) crucially depends on the i.i.d normality of the data generating process. To understand how departures from this requirement affect the null distribution of \widehat{S} , suppose the process $\{Y_t\}$ is Gaussian but serially correlated. Conditional on high past values of Y_t , if the process is positively correlated, future values of Y_t are more likely to be also high. The nominal size of the test is understated and it is likely to detect spurious positive skewness, resulting simply from serial dependence. Lomnicki (1961) shows that if $\{Y_t\}$ admits a moving-average representation, then $\sqrt{T}(\widehat{S} - S)$ is asymptotically normal with zero mean and variance $6 \sum_{j=-\infty}^{+\infty} \rho_j^3$, where ρ_j is the autocorrelation of order j of Y_t .

When data is i.i.d. but non-Gaussian, the asymptotic variance of (3.9) is also incorrect. In particular, if the sample comes from a symmetric but leptokurtic distribution, a test based on (3.9) would also reject the null of symmetry too often.

In this case, Stuart and Ord (1987) shows that $\sqrt{T}(\widehat{S} - S)$ is asymptotically normal with zero mean and variance

$$9 - 6m_4/m_2^2 + m_6/m_2^3, \quad (3.10)$$

where m_4/m_2^2 is the coefficient of kurtosis, or the standardized fourth moment. This asymptotic variance requires the existence of the sixth moment, which may be too restrictive in the context of financial returns, as documented by Jansen and de Vries (1991), Loretan and Phillips (1994) and Lima (1997). More recently, Bai and Ng (2005) derived the limiting distribution of \widehat{S} in the general case of non-Gaussian and (weakly) dependent data. However, their skewness test applies only to series that are stationary up to the sixth-order, which is not the case of many financial returns that typically exhibit fat tails.

Lastly, despite being consistent, \widehat{S} is biased in finite samples, simply because sample moments are not necessarily unbiased estimators of their population counterparts. Specifically, if \widehat{m}_k is the k -th ordered sample moment, then $\mathbb{E}[\widehat{m}_2] = m_2(T-1)/T$ and $\mathbb{E}[\widehat{m}_3] = m_3(T-1)(T-2)/T^2$. There are simple small sample corrections for \widehat{S} , but they are not unique and different statistical packages usually employ different formulations with different finite-sample properties, as discussed by Joanes and Gill (1998). However, when the sample size is large this bias is negligible.

It is also known the estimation of higher-order moments is sensitive to outliers, and the presence of a few extreme returns may lead an statistical test to indicate asymmetry in an otherwise symmetric distribution. In order to mitigate this effect robust measures of skewness have been suggested. Analogous to the role that the median and the interquartile range play as robust measures of location and scale, Hinkley (1975) considers a general class of asymmetry measures based on quantiles,

$$S_R(p) = \frac{(Q_{1-p} - Q_{0.5}) - (Q_{0.5} - Q_p)}{Q_{1-p} - Q_p},$$

where Q_p for $p \in (0, 1/2)$ denotes the p -th quantile, i.e. $F(Q_p) = p$ and $Q_{0.5}$ is the median. Bowley (1937) skewness statistic is a particular case in which quartiles

($p = 0.25$) are used, and

$$\sqrt{T} \widehat{S}_R(0.25) \xrightarrow{d} \mathcal{N}(0, 1.839) \quad (3.11)$$

in i.i.d. normal samples under the null of symmetry. However, the choice of $p = 0.125$ is also a valid statistic based on octiles and puts more weight on the tails of the distribution than Bowley's measure. It is not clear which p should be chosen for each application, and one could also integrate $S_R(p)$ over p to obtain $(\mathbb{E}Y - Q_{0.5})/\mathbb{E}|Y - Q_{0.5}|$, a skewness measure free of the p parameter and bounded on $(-1, 1)$. In spite of being robust to outliers, the limiting distributions of the above statistics, to the best of our knowledge, were derived only for the i.i.d. normal or exponential cases, by Moors et al (1996) and for the general i.i.d. case, by Ngatchou-Wandji (2006).

Neither (3.9), (3.11) or the Bai and Ng (2005) tests are adequate to detect asymmetry in the empirical distribution of stock returns, which has been shown to exhibit leptokurtosis, serial dependence and time-varying volatility, specially when calculated at short (daily or weekly) intervals (Tsay, 2005). Also, the possibility of arbitrage suggests that asymmetry, at least unconditionally, is not likely to be very substantial or persistent. Singleton and Wingender (1986) confirmed earlier evidence that *ex-post* common stock returns tend to be positively skewed in the cross-section, but showed that this phenomenon is not persistent for individual stocks or portfolios. Hence, better measures are required and in order to be valid and useful for financial data, symmetry tests must not only be robust to the stylized facts of stock returns but should also be powerful enough to detected small but significant departures from symmetry. Both of these desirable features are usually in conflict with one another, so it seems natural and appropriate to apply more than a single symmetry test as they vary in terms of robustness and power.

In the last few years a handful of symmetry tests were specially designed to be applied to financial data, and four of these tests are appropriate to our study. Lisi (2007) uses the traditional skewness coefficient (3.6) applied to the residuals of a properly fitted ARMA model of returns, relying on bootstrapped standard errors for inference. The method of Chen and Lin (2008) explores the fact that for any process Z_t that is symmetric, mean-zero and stationary, $\mathbb{E}[\phi(Z_t)] = 0$ for

any odd function $\phi(\cdot)$. The test of Bai and Ng (2001) for conditional symmetry is based on the empirical distribution of the innovations of a fully specified ARMA-GARCH model of returns. Finally, the Maasoumi and Racine (2009) test, which measures the distance between kernel density estimates of Y_t and its rotated series $\tilde{Y}_t \equiv -Y_t + 2\mu_y$, where μ_y is the location parameter of Y_t . Under symmetry, Y_t and \tilde{Y}_t are equally distributed, so the distance between kernel estimates of their densities must be close to zero. The test can also be applied to consistent estimates of the innovations ϵ_t and $\tilde{\epsilon}_t$, serving as a test of both conditional and unconditional symmetry of returns.

In order to perform these tests, we must specify and estimate a model such as (3.5) for each stock in our sample. We used an ARMA process for the mean equation and a APARCH specification for the variance equation, following Ding, Granger and Engle (1993),

$$h_t^\delta = \omega + \sum_{i=1}^p \alpha_i (|\eta_{t-i}| - \gamma_i \eta_{t-i})^\delta + \sum_{i=1}^q \beta_i h_{t-i}^\delta. \quad (3.12)$$

Several GARCH models are nested in (3.12), and in practice we set $p = q = 1$ and $\delta = 2$ for most stocks in our sample, which effectively reduces (3.12) to the GJR-GARCH(1,1) model of Glosten, Jagannathan and Runkle (1993). The APARCH model is a fairly general and yet parsimonious specification to capture the stylized facts of conditional variance in stocks returns (Teräsvirta, 2009). Most importantly, we assume that ϵ_t follows a skewed version of the Student-t distribution, in the sense of Fernandez and Steel (1998), which is a flexible specification that accommodates asymmetry and leptokurtosis in the data. Specifically, for each stock $i = 1, \dots, N$ in our sample, we assume

$$\epsilon_{i,t} \sim \text{skewed-t}(\xi_i, \nu_i), \quad (3.13)$$

where parameters ν_i and ξ_i control the tail mass and asymmetry, respectively (see Appendix C for further details on this distribution). This flexibility allows for consistent estimation of θ and λ under the possibility of ϵ_t being truly skewed. Newey and Steigerwald (1997) showed that if the conditional mean is not identically zero, the quasi-maximum likelihood estimator (QMLE) of conditional heteroscedasticity models such as (3.5) and (3.12) is consistent if both the true and assumed error

density are symmetric. When this symmetry condition does not hold, either a flexible distribution (i.e., that allows for skewness) must be assumed for ϵ_t or an additional parameter must be included in the mean equation so that the possibly non-zero location of the disturbance process can be identified. By considering a flexible distribution and estimating the mean and variance equations jointly, we preserve consistency of the QMLE of (3.5) and (3.12) under the possibility that the true innovations are skewed.

3.2.1 Unconditional Asymmetry Tests

Lisi (2007) addresses the problem of correctly quantifying the variance of \widehat{S} by bootstrapping the residuals of a properly fitted ARMA model of returns. Using the notation in (3.5), a skewness coefficient is calculated for the filtered series, $\widehat{S}(\eta_t)$, and the variance of this statistic under the null of symmetry is obtained by resampling “symmetrized” residuals: if $\{\eta_1, \dots, \eta_T\}$ is the original sample of residuals and $\eta_t^* \equiv |\eta_t - me(\eta_t)|$, where $me(\eta_t)$ is the median of η_t , then

$$\widetilde{\eta}_t = me(\eta_t) + z_t \eta_t^*, \quad t = 1, \dots, T$$

where z_t is a Rademacher random variable, i.e., $P(z_t = -1) = P(z_t = +1) = 1/2$, is a reallocation of the original residual series that imposes symmetry. Resampling with replacement from $\{\eta_1, \dots, \eta_T\}$ and calculating $\widehat{S}(\widetilde{\eta}_t)$ for each resampled series yield a bootstrapped distribution of $\widehat{S}(\eta_t)$ under the null of symmetry. This method does not exactly test the unconditional skewness of y_t but is an approximation that may work well when the conditional mean does not vary much, which seems to be the case for short-term stock returns. Lisi provides a simulation study showing that inference with the suggested resampling method has better size and power than with the asymptotic variance (3.10), with the advantage of requiring existence of the third moment of y_t only.

Chen and Lin (2008) proposed a family of symmetry tests for the unconditional asymmetry of a return series y_t that is robust to serial dependence and leptokurtosis. The test explore the fact that $\mathbb{E}[\phi(X)] \neq 0$ implies asymmetry of X for any odd function ϕ . For three different choices of odd functions, namely

$\phi_s(z) = z^3$, $\phi_c(z) = z/(1 + z^2)$ and $\phi_p(z) = \arctan(z)$, the authors develop two types of tests: the H test, based on the statistic

$$H_T = \frac{1}{T} \left[\sum_{t=1}^T \phi(z_t) \right]^\top \widehat{\Omega}_T^{-1} \left[\sum_{t=1}^T \phi(z_t) \right], \quad (3.14)$$

where $\widehat{\Omega}_T^{-1}$ is a HAC estimator of the long-run variance of $T^{-1/2} \sum_{t=1}^T \phi(z_t)$ and $z_t = (y_t - \bar{y})/\widehat{\sigma}_y$ is the standardized return; and the K test, based on the statistic

$$K_T = \frac{1}{T} \left[\sum_{t=1}^T \phi(z_t) \right]^\top \widehat{C}_T^{-1} \left[\sum_{t=1}^T \phi(z_t) \right], \quad (3.15)$$

in which the long-run variance estimate $\widehat{\Omega}_T$ is replaced by a weighting matrix \widehat{C}_T (see their paper for details). While $H_T \xrightarrow{d} \chi^2(1)$, the distribution of K_T is non-standard and its quantiles must be obtained by simulation: the 90% (95%) critical value is 28.31 (45.40). Applying each odd function ϕ_s , ϕ_c and ϕ_p to (3.14) or (3.15) result in six different test statistics: HS , HC , HP , KS , KC and KP . Note that in spite of being robust to the mentioned problems, these tests will only detect departures from symmetry and cannot discriminate against positive or negative asymmetry.

Maasoumi and Racine (2009) developed an entropy-based test of symmetry of a stationary (continuous or discrete) process $\{X_t\}$ based on the distance between the kernel density estimates \widehat{f}_1 of X_t and \widehat{f}_2 of $\widetilde{X}_t \equiv 2\mathbb{E}[X_t] - X_t$, i.e., a rotation of X_t about its mean. If X_t is symmetric, then $f_1(x) = f_2(\widetilde{x})$ almost surely, and the proposed test statistic \widehat{S}_ρ is the integrated squared difference among the square-root of densities,

$$\widehat{S}_\rho = \frac{1}{2} \int_{-\infty}^{+\infty} [\widehat{f}_1^{1/2}(x) - \widehat{f}_2^{1/2}(\widetilde{x})]^2 dx. \quad (3.16)$$

Instead of developing asymptotic critical regions, the null distribution of \widehat{S}_ρ is obtained by bootstrap: while the asymptotic critical values do not depend on the bandwidth of the kernel function (as this is a quantity that vanishes asymptotically), the value of the test statistic depends directly on the bandwidth. This is a major drawback for an asymptotic-based test, as variability of the test statistic

across bandwidths may be substantial. The empirical distribution of \widehat{S}_ρ under the null of symmetry is obtained by resampling with replacement from a pooled sample $Z = \{X_1, \dots, X_T, \widetilde{X}_1, \dots, \widetilde{X}_T\}$ of size $2T$, and calculating \widehat{S}_ρ^* for each bootstrapped sample Z^* . With sufficiently many replications, say $\widehat{S}_{\rho,1}^*, \dots, \widehat{S}_{\rho,B}^*$, the p -value of the test statistic is the proportion of bootstrapped statistics \widehat{S}_ρ^* that are at least as extreme as the value obtained from the original sample, $\sum_{b=1}^B \mathbf{1}\{\widehat{S}_{\rho,b}^* \geq \widehat{S}_\rho\}/B$. As suggested by the authors, we used $B = 399$ replications and the stationary bootstrap of Politis and Romano (1994). Since the test statistic is a distance measure, it is always non-negative and can not distinguish between right or left asymmetry as well.

3.2.2 Conditional Asymmetry Tests

As the model (3.5), (3.12) and (3.13) is estimated for each stock $i = 1, \dots, N$, a direct test of conditional asymmetry can be based on the statistical significance of the estimates $\widehat{\xi}_i$ of the innovation distribution. As discussed previously, the QML estimation under (3.13) is consistent and asymptotically normal, providing the distribution necessary for significance testing.

Bai and Ng (2001) developed a test of symmetry for the residuals ϵ_t for a model such as (3.5) that does not require stationary nor i.i.d. data. Their test is consistent, asymptotically distribution-free and requires only consistent estimates of the residuals, i.e., of the mean and variance equations of (3.5) in the sense of Newey and Steigerwald (1997). Their test is based on the empirical distribution functions of ϵ_t and $-\epsilon_t$, which under the null of symmetry have the same distribution, denoted here by $F(\epsilon)$. The building blocks of their CS statistic are as follows. Let

$$W_T(x) = \frac{1}{\sqrt{T}} \sum_{t=1}^T [\mathbf{1}\{\epsilon_t \leq x\} - \mathbf{1}\{-\epsilon_t \leq x\}] \quad (3.17)$$

denote the difference between the number of ϵ_t and the number of $-\epsilon_t$ that are less than or equal to x , then divided by \sqrt{T} . Under symmetry of ϵ_t , $W_T(x)$ should be small at all values of x , and if ϵ_t was observed for all t , then $\max_x |W_T(x)|$ would be a natural candidate for testing conditional symmetry. Since $\widehat{\epsilon}_t$ must

be used, let \widehat{W}_T denote the feasible estimator of (3.17) and for $x \leq 0$ define $S_T(x) = \widehat{W}_T(x) - \widehat{W}_T(0) + \int_x^0 h_T^-(y)dy$, where

$$h_T^-(y) = g_T(y)f_T(y) \left[\int_{-\infty}^y g_T(z)^2 f_T(z) dz \right]^{-1} \int_{-\infty}^y g_T(z) d\widehat{W}_T(z).$$

For $x > 0$, define $S_T(x) = \widehat{W}_T(x) - \widehat{W}_T(0) - \int_0^x h_T^+(y)dy$, where

$$h_T^+(y) = g_T(y)f_T(y) \left[\int_y^{+\infty} g_T(z)^2 f_T(z) dz \right]^{-1} \int_y^{+\infty} g_T(z) d\widehat{W}_T(z)$$

and where f_T is a kernel estimator of the density $f(\epsilon_t)$ based on a sample of size T , and g_T is an estimator of $g = \dot{f}/f$. The proposed test is based on the statistic

$$CS = \max_x |S_T(x)| \xrightarrow{d} \max_{0 \leq s \leq 1} |B(s)|, \quad (3.18)$$

where $B(r)$ is a standard Brownian motion. The asymptotic critical values are obtained by simulation, and are 2.21 and 1.91 for the 5% and 10% significance levels, respectively.

Lastly, the Maasoumi and Racine \widehat{S}_ρ test (3.16) described in section 3.2.1 can also be applied to the residuals ϵ_t for a test of conditional symmetry. The only necessary adjustment is on the resampling methodology for deriving the null distribution of the test statistic, which in this context can be the simple i.i.d. bootstrap.

3.2.3 Empirical Results of Asymmetry Tests

Our sample consists of daily (log) returns of Brazilian stocks traded at the Bovespa from July 1994 to July 2009, covering the first 15 years of low inflation after the Real stabilization plan. Excess returns are calculated over the 1-day interbank deposit certificate (CDI), a common proxy for the risk-free rate in Brazil. Only companies whose most liquid stock were negotiated in at least 70% of the trading sessions were selected. Less than 40 companies satisfy this criteria during the entire 15 years, so our sample would be extremely small and biased towards the survivors if selection were restricted to the full sample period. In order to mitigate this problem we divided the 1994-2009 period into five non-overlapping three year

sub-periods, and for each sub-period we picked every company that satisfied the 70% presence criteria¹. This subsampling strategy also helps to detect temporary departures of symmetry that might occur. We collected a total of 481 series of returns across all subperiods, with some firms contributing to more than one series.

We first applied the traditional and robust tests (3.9) and (3.11), and report in Table 3.1 the number of firms with significant skewness at a 10% nominal level, by sub-period and regulatory status. As discussed, these tests are likely to have an incorrect size when applied to serially correlated and heavy-tailed data. The traditional, i.i.d. normal test for skewness rejects symmetry in almost 70% of cases, while the robust (quantile-based) test, also based on i.i.d. normal returns, rejects less than 30% of cases. Both measures however indicate a decrease in the relative number of asymmetric cases in more recent periods and a predominance of positive asymmetry.

The results of the conditional and unconditional tests described in Sections 3.2.1 and 3.2.2 applied to the same sample of stock returns are reproduced in Tables 3.2, 3.3 and 3.4, for a significance levels of 5% and 10%. Out of the 481 return series considered, from 15% to 30% were considered unconditionally asymmetric at the 5% level by the various tests. Under a 10% significance, from 23% to 41% of cases are considered asymmetric, depending on the specific test, which shows how oversized the traditional test (3.9) can be when applied to non-normal, dependent data. Lisi's test, the only unconditional test capable of detecting left and right asymmetry, also shows that the vast majority of asymmetric returns are positively skewed. In conditional terms, asymmetry is also predominantly positive and appears to be more pervasive across firms and periods, as some 23% to 48% (at a 5% level) or 34% to 56% (at a 10% level) of innovation processes were rejected by the $\hat{\xi}$, \hat{S}_ρ and CS tests.

¹Empirical studies with Brazilian stocks usually have to deal with the issue of small (but growing) liquidity. There are not too many companies quoted in every trading session and, to make matters worse, most firms have more than one equity class outstanding (usually two classes, one with and another without voting rights). The trade-off in lowering the liquidity threshold to increase the sample size (cross-sectional units) is the introduction of the statistical problems associated with non-synchronous or infrequent trading and missing data. We found that requiring presence in at least 70% of trading sessions generates samples with at least 100 different companies every year, each with no more than 5 consecutive days of missing data.

3.3 The Relation Between Asymmetry and Regulatory Status

In this section we pursue four different approaches to test whether there is significant association among return asymmetry S and regulatory status R among the stocks in our sample, i.e., a test of hypothesis (3.1). It would be inappropriate to impose any *a priori* structure on $f_{S,R}$, the very object we want to study, so all methods employed to test whether S and R are related are nonparametric. Gibbons and Chakraborti (2003) and Li and Racine (2007) provide an extensive discussion of the tests considered in this section.

3.3.1 Contingency Tables

Hypothesis (3.1) states independence among return asymmetry S , a continuous variable not directly observable, and regulatory status R , a dichotomous variable indicating whether the firm is regulated or not. If S could be observed and discretized in some economically meaningful way, a contingency table could be formed and the relation between S and R could be tested. In particular, discretizing S into two categories would allow Fisher's exact test to be carried. For concreteness, let there be n_0 unregulated and n_1 regulated firms in our sample, and assume that the true asymmetry s_i of stock returns of firm $i = 1, \dots, n_0 + n_1$ is observed. Let $s^* > 0$ be some cutoff point (determined in terms of economic significance, for instance) such that we classify as "symmetric" every return distribution i such that $|s_i| < s^*$, and "asymmetric" otherwise. Let k denote the total number of firms categorized as asymmetric, and among these k firms let x be the subset of regulated ones. These four values completely describe a contingency table such as Table 3.5.

Under the null hypothesis of independence and with fixed marginals n_0 , n_1 and k , Fisher (1922) derived the exact probability of observing x in the upper left cell,

$$P_H(x | n_0, n_1, k) = \binom{n_1}{x} \binom{n_0}{k-x} / \binom{n_0 + n_1}{k} \quad (3.19)$$

which is the hypergeometric distribution. For 2×2 tables, Fisher's exact test for independence is equivalent to test whether the odds ratio,

$$OR(x) = \frac{x}{n_1 - x} \bigg/ \frac{k - x}{n_0 - (k - x)}, \quad (3.20)$$

is significantly different from 1, a two-sided hypothesis. Our conjecture is that a larger proportion of regulated companies tend to exhibit asymmetry, so the appropriate alternative is $OR > 1$, a one-sided test. The exact p -value of this one-sided test is obtained by adding the probabilities of each table arrangement that yield an odds ratio higher than (3.20), conditionally on the observed marginals. The only way to increase OR keeping the marginals fixed is to transfer observations from the cell (Regulated \cap Symmetric) to (Regulated \cap Asymmetric) and, at the same time, from (Unregulated \cap Asymmetric) to (Unregulated \cap Symmetric), i.e., by increasing x . This can be done until any of the "donating" cells becomes empty. Then, the one-sided p -value of Table 3.5 for testing $OR(x) > 1$ is

$$p\text{-val} = \sum_{i=1}^L P_H(x + i | n_0, n_1, k), \quad (3.21)$$

where $L = \min\{n_1 - x, k - x\}$.

In practice however S must be estimated, and the sampling error in \widehat{S} does not allow us to observe the true entries in Table 3.5. Fisher's test applied on the actually observed table would inform an incorrect p -value due to possible misclassification. Fortunately, under some assumptions, we can work around this problem. Consider any estimator \widehat{S} of S and a test of $S = 0$ with rejection probability α and power $1 - \beta$. We expect the following misclassification: $\alpha\%$ of the companies with symmetric returns will fall in the asymmetric bin, while $\beta\%$ of the companies with asymmetric returns will be mistaken for symmetric ones. If the test error probabilities are independent of R , and if the power of the test is constant under the alternative ($S \neq 0$), the cross tabulation actually observed would be given by the entries A , B , C and D shown in Table 3.6, in which actual entries are rounded to the nearest integer.

Row totals $n_0 = A + B$ and $n_1 = C + D$ would still be observed without error. Given α and β , the correct p -value for the one-sided test $OR > 1$ under

measurement error in S can be calculated by

$$p\text{-val}(\alpha, \beta) = \sum_{i=1}^{L(\alpha, \beta)} P_H \left(\frac{A - n_1 \alpha}{1 - (\alpha + \beta)} + i \mid n_0, n_1, \frac{A + C - (n_0 + n_1) \alpha}{1 - (\alpha + \beta)} \right), \quad (3.22)$$

where $L(\alpha, \beta) = \min \{(n_1(1 - \beta) - A), (C - n_0 \alpha)\} / (1 - (\alpha + \beta))$.

For each symmetry test reviewed in Sections 3.2.1 and 3.2.2, we calculate (3.22) with $\alpha = 5\%$ and 10% and $\beta = 5\%$, 20% and 35% to allow for different sizes and powers these tests may have in our case. Tables 3.7, 3.8 and 3.9 show the results of this sensitivity analysis. Each asymmetry test generates a contingency table such as Table 3.6 with an associated odds ratio. In only a few cases there was evidence that the true odds ratio was greater than 1, specially during the first period (1994-1997) and for conditional asymmetry. However, the evidence for an odds ratio different than one (a two-sided test) was much stronger. Most of the cases in which the null was rejected the estimated odds ratio was less than one, which means that regulated stocks exhibited *less* skewed returns than their non-regulated counterparts. This phenomenon was significant in all 5 subperiods, specially regarding conditional asymmetry. Table 3.10, which aggregates and summarizes the results presented in Tables 3.7, 3.8 and 3.9, provides supporting evidence, which seems to be robust to different values of α and β .

A limitation of this approach is that it cannot be applied to every observed table. For any given pair of test error probabilities (α, β) , there is a lower and an upper bound for the observed odds of being asymmetric (A/B or C/D), given by

$$\alpha / (1 - \alpha) \leq A/B \text{ or } C/D \leq (1 - \beta) / \beta, \quad (3.23)$$

beyond which the proposed transformation can not be applied as it would generate negative “true” entries. This is so because we are assuming that the symmetry tests have constant power across the entire region of asymmetry. As test power usually increases as we move further away from the null, and (3.22) does not take this into account, a fixed level of β tends to overestimate (underestimate) the true power of the asymmetry test for stocks that are closer to (further from) the null. In tables with extreme odds, beyond the bounds in (3.23), the transformation proposed fails to generate sensible values for the associated true classification.

3.3.2 Equality of the Conditional Densities

In this section we perform three types of tests to verify the validity of hypothesis (3.2). The first is based on the ranks of the R variable when the sample is ordered by the estimated asymmetry coefficients. The other two are based on estimates of the conditional distribution in (3.2). Non-smooth estimates of $F_{S|R}(s|r=0)$ and $F_{S|R}(s|r=1)$ are compared with a modified version of the traditional Kolmogorov-Smirnov, Cramr-von Mises and Anderson-Darling tests, while smooth (kernel-based) estimates are compared with the entropy metric of Maasoumi and Racine (2002).

Consider the results of an asymmetry statistic \widehat{S} applied to a set of N stock returns, and sort the bivariate sample $\{(\widehat{s}_i, r_i), i = 1, \dots, N\}$ according to the value of \widehat{s} . The resulting arrangement of R in the ordered sample

$$\{(\widehat{s}_{(1)}, r_{(1)}), \dots, (\widehat{s}_{(N)}, r_{(N)})\}$$

can be tested for randomness, which is also a test of (3.2). Too few runs in the sequence of zeros and ones in $\{r_{(1)}, \dots, r_{(N)}\}$ is evidence against the equality of the above conditional distributions. One advantage of this test is that it is robust to small perturbations in the values of \widehat{s}_i , as long as these perturbations do not change the ranks too much, so the problem of measurement error in S is alleviated.

Table 3.11 shows the results of the runs test applied to each sequence of $\{r_{(i)}\}$ formed after ordering the estimates from the 11 tests of asymmetry considered. There is not enough evidence to reject (3.2) on the basis of this test: only for the 1997-2000 period there exists some indication of non-randomness in unconditional asymmetry, as 5 out of the 8 tests rejected the null.

Another strategy for testing (3.2) is to partition the sample $\{(\widehat{s}_i, r_i), i = 1, \dots, N\}$ conditional on the value of R and to apply one of the classical tests for equality of distribution in the two samples

$$\{(\widehat{s}_1, 0), \dots, (\widehat{s}_{n_0}, 0)\} \text{ and } \{(\widehat{s}_1, 1), \dots, (\widehat{s}_{n_1}, 1)\}$$

obtained. We use more powerful versions of the classical tests of Kolmogorov-Smirnov (K-S), Cramr-von Mises (C-vM) and Anderson-Darling (A-D) proposed

by Zhang (2006). Concretely, let \widehat{F}_0 , \widehat{F}_1 and \widehat{F} denote the empirical distribution function of the unregulated, regulated and combined samples, respectively. Instead of the usual chi-squared distance criteria, Zhang (2006) suggests the following likelihood-ratio statistic

$$G(s) = 2 \sum_{j=0}^1 n_j \left[\widehat{F}_j(s) \log \frac{\widehat{F}_j(s)}{\widehat{F}(s)} + (1 - \widehat{F}_j(s)) \log \frac{1 - \widehat{F}_j(s)}{1 - \widehat{F}(s)} \right] \quad (3.24)$$

as the distance metric for the K-S, C-vM and A-D tests, where n_0 and n_1 are the number of unregulated and regulated companies in the sample. The likelihood-ratio version of these statistics are, respectively,

$$Z_K = \sup_s |G(s)| \quad (3.25a)$$

$$Z_C = \int_{-\infty}^{\infty} G(s) d\widehat{F}(s) \quad (3.25b)$$

$$Z_A = \int_{-\infty}^{\infty} G(s) \widehat{F}^{-1}(s) [1 - \widehat{F}(s)]^{-1} d\widehat{F}(s) \quad (3.25c)$$

and their null distributions are obtained by enumerating all the possible and equally likely $n!/(n_0!n_1!)$ values of the statistic. A simulation exercise provided by Zhang (2006) showed that the the tests (3.25) are at least as powerful as their original versions for location differences, and much more powerful for changes in scale or shape. For our sample of stocks, obtaining an exact p -value is infeasible as $n_0 + n_1 \approx 100$, so an approximate p -value was obtained by Monte Carlo, randomly selecting 10,000 values from the null distribution.

Table 3.12 reports the results of the Z_A test, as the Anderson-Darling is generally considered the most powerful among the three classical tests. The results obtained with Z_K and Z_C were similar and are available from the authors by request. The general conclusion is analogous to the results of the runs test: only during the 1997-2000 period there seemed to exist a significant difference in the distribution of asymmetry among regulated and non-regulated stocks, both in conditional and unconditional terms. While the evidence from the Z_A test for that period is strong, it does not indicate the source of discrepancy between $\widehat{F}_{S|R}(s|r=0)$ and $\widehat{F}_{S|R}(s|r=1)$, i.e., whether it is due to location, scale or shape differences.

Using the same partitioned sample it is possible to estimate the conditional density functions by kernel methods, calculate some metric of discrepancy among the estimated densities and decide whether the calculated distance is statistically significant. We performed a test based on kernel smoothing methods proposed by Maasoumi and Racine (2002) and Granger, Maasoumi and Racine (2004). Their distance metric, applied to our context, can be written as

$$\widehat{D} = \int_{-\infty}^{\infty} \left[\widehat{f}_{R|S}^{1/2}(s|r=0) - \widehat{f}_{R|S}^{1/2}(s|r=1) \right]^2 ds, \quad (3.26)$$

and the null distribution of \widehat{D} is obtained by bootstrapping the statistic from the combined sample $\{\widehat{s}_1, \dots, \widehat{s}_{n_0}, \widehat{s}_{n_0+1}, \dots, \widehat{s}_{n_0+n_1}\}$. Table 3.13 contains the results of the \widehat{D} test, and once again the period from 1997 to 2000 stands out as the only one with sufficient evidence of inequality between the conditional densities, for both unconditional and conditional asymmetry. The \widehat{D} test is also incapable of providing a precise source for the discrepancy between the conditional densities.

In conclusion, we were not able to reject (3.2) consistently (i.e., for the majority of the asymmetry criteria and across all the sampled periods) by any of the three different techniques employed to measure the degree of discrepancy of the density of S conditional on the different regulatory statuses.

3.3.3 Location and Scale Differences

In Sections 3.3.1 to 3.3.2 we failed to reject the hypothesis (3.1) and (3.2) against the general alternatives H_A and H'_A , except for the period from 1997 to 2000. As we only employed robust, nonparametric techniques to test such hypothesis, this could be simply the result of the relative lower power these tests have when compared to their parametric counterparts. One possible way to improve inference is to develop more specific alternative hypothesis, motivated by the theoretical implications of our conjecture, which naturally lead to more powerful tests.

One expected outcome of our conjecture is that for any asymmetry measure $\mathcal{S} \in \mathbb{R}^+$ that does not discriminate among left or right skewness, the distribution of \mathcal{S} among regulated stocks must have a larger location parameter than the dis-

tribution of \mathcal{S} among unregulated firms. The Wilcoxon Rank-Sum test can be used to test whether there is a significant discrepancy among the locations of these distributions, in the form of hypothesis (3.4). The test statistic is $W = \sum_{i=1}^N i r_{(i)}$, where the sequence of 0's and 1's in $\{r_{(i)}\}$ is sorted according to the respective asymmetry estimates \hat{s}_i as used for the runs test in Section 3.3.2. While an exact null distribution can be calculated by enumeration when N is small, the normal approximation is accurate enough for combined samples of at least 12 observations, which is our case.

Table 3.14 summarizes the results of the Wilcoxon Rank-Sum test of (3.4), for the various asymmetry concepts and sampled periods. The general conclusion from this test is somewhat similar and confirms the results from the previous tests: the period from 1997 to 2000 stands out as the only one in which evidence for a difference among the conditional densities is strong enough, and, as indicated by previous test, in the sense that returns of regulated stocks are less asymmetric on average. Table 3.14 reports one-sided p -values based on the alternative hypothesis that the location parameter of $f_{\mathcal{S}|R}(s|r = 1)$ is larger than the location parameter of $f_{\mathcal{S}|R}(s|r = 0)$, i.e., $\theta \geq 0$. However, for the 1997-2000 period, the Wilcoxon test applied to 7 out of the 8 unconditional asymmetry measures reported p -values larger than 0.95, which means that the location parameter for the regulated group is smaller and statistically significant at 5% in a one-sided test in the opposite direction ($\theta \leq 0$). The same conclusion can be drawn regarding conditional asymmetry, as 2 out of the 3 metrics yield similar results for the 1997-2000 period.

One additional implication of our conjecture is that the distribution of any metric S that also informs the sign or direction of asymmetry is likely to have a larger variance among regulated companies. This formally translates to the hypothesis (3.3), which can be non-parametrically tested by the Ansari-Bradley or Mood statistics if we use the skewness metric of Lisi, for unconditional symmetry, or the ξ_i parameter of the GARCH innovation in (3.13), for unconditional symmetry². Table 3.15 reports the results of these tests, and shows that the larger

²It is statistically valid to test for scale discrepancies after the location alternative has been tested, as these tests are asymptotically independent when the set of weights of the linear rank test statistic obeys some usual conditions, namely symmetry about their mean. See Anderson

variance effect might have happened during the 1994-1997 period in terms of unconditional asymmetry. The Mood test indicates the same also happened from 2000 to 2003. And regarding the 1997-2000 period, we refrain from drawing any further conclusions, as the location parameters of $f_{S|R}(s|r = 0)$ and $f_{S|R}(s|r = 1)$ are likely to be different (and unknown) and the distribution of the scale tests is not well defined in this case.

3.4 Conclusions

We conjectured that regulated companies are subject to certain phenomena absent in a unregulated environment, potentially affecting the symmetry stock returns. Using a sample of Brazilian publicly traded firms, we could not find consistent differences in the distribution of cross-sectional asymmetry coefficients among regulated and non-regulated firms, i.e., across several periods and different measures of asymmetry. There is evidence, however, that during the period from 1997 to 2000, non-regulated companies exhibited *less* asymmetric returns.

The period from 1997 to 2000 marks the emergence and early development stages of economic regulation of key industries in Brazil, and their transition from state-owned monopolies to private, regulated businesses. The federal agencies in charge of the electricity, telecommunications and oil sectors, ANEEL, ANATEL and ANP respectively, were created in the period from 1996 to 1998, in conjunction with other major reforms such as privatization and liberalization. Other industries such as gas, water and transportation also witnessed the appearance of state level regulation during the same period. It is likely that such significant changes in the environment of regulated businesses had an impact on the empirical distribution of stock returns in that period. Even under such radical transformations, we found evidence that regulated firms exhibited less skewed returns than their unregulated counterparts. The mechanism by which the establishment of large regulatory bodies affect the shape of ex-post stock returns is unclear and would be an interesting topic for further investigation.

(2001) for a proof.

The empirical results in this paper are subject to limitations of practical nature. First, due to the size of the market, the number of industries surveyed is relatively small and might not allow for sufficient variation in return asymmetry (which, after all, is not likely to be substantial or persistent) to appear. Second, we had to deal with the fact that asymmetry we cannot be observed but must be estimated. There is measurement error in the very quantity of interest, and we tried to mitigate this problem by using only robust nonparametric tests, specially those based on ranks, which are unlikely to vary much as a result of sampling error. Another approach to alleviate this problem is to conduct a Monte Carlo sensitivity analysis on the results of each test performed in Section 3.3. By randomly drawing a new, cross-sectional set of estimates from the sampling distribution of the asymmetry statistics, each test performed in section 3.3 could be repeated a large number of times, and the relative impact of sampling error could be assessed. This additional exercise would be important if we were able to reject the null hypotheses (3.1) and (3.2) with the primary set of estimated asymmetry coefficients, as a robustness check. However, since the data were not able to generate sufficient evidence against the null, this additional effort seems unnecessary.

This chapter is based on *Economic Regulation and Asymmetry in Ex-Post Stock Returns*, joint work with Regio Martins.

3.5 Tables and Figures

Table 3.1: Firms with significant skewness by the (3.9) and (3.11) tests.

Period	Type of Firm	Total	Skewness \widehat{S}		Robust $S_R(0.25)$	
			Positive	Negative	Positive	Negative
1994-1997	Regulated	28	22	0	12	0
	Unregulated	59	45	2	10	10
1997-2000	Regulated	29	10	13	4	1
	Unregulated	50	35	5	9	8
2000-2003	Regulated	37	13	5	1	4
	Unregulated	50	30	5	11	3
2003-2006	Regulated	42	29	1	11	2
	Unregulated	54	32	1	21	1
2006-2009	Regulated	47	22	6	6	4
	Unregulated	85	35	21	11	8
Total		481	273	59	96	41

Note: Entries are the number of companies with statistically significant \widehat{S} and $\widehat{S}_R(0.25)$, assuming i.i.d. normal returns and a 10% significance level.

Table 3.2: Results of the asymmetry tests described in Section 3.2.1, at 5% nominal size

Period	Sampled Firms		Unconditional Tests										\widehat{S}_ρ
	Type	Quantity	$\widehat{S}(+)$	$\widehat{S}(-)$	HS	HC	HP	KS	KC	KP			
1994-1997	Regulated	28	8	0	3	13	8	15	15	15	12		
	Unregulated	59	26	0	17	30	30	28	30	33	23		
1997-2000	Regulated	29	3	2	1	2	2	2	1	1	8		
	Unregulated	50	18	0	10	16	15	5	5	9	20		
2000-2003	Regulated	37	9	0	6	7	7	10	8	9	7		
	Unregulated	50	15	1	8	18	16	15	16	15	13		
2003-2006	Regulated	42	16	0	6	17	18	11	16	16	12		
	Unregulated	54	22	0	13	22	22	17	18	17	19		
2006-2009	Regulated	47	3	0	2	6	4	3	5	4	5		
	Unregulated	85	14	2	8	16	16	8	8	8	11		
Total		481	134	4	74	147	138	114	122	127	130		

Note: Entries are the number of firms with asymmetric returns according to the tests proposed by Lisi (the skewness coefficient \widehat{S} with bootstrapped standard errors), Chen and Lin (the H and K tests), and Maasoumi and Racine (the \widehat{S}_ρ test), all at a 5% level of significance. In this table only the \widehat{S} test is capable of distinguishing positive (right-skewed) from negative (left-skewed) asymmetry. Its results are reported with indicative (+) and (-) signs.

Table 3.3: Results of the asymmetry tests described in Section 3.2.1, at 10% nominal size

Period	Sampled Firms		Unconditional Tests										\widehat{S}_ρ
	Type	Quantity	$\widehat{S}(+)$	$\widehat{S}(-)$	HS	HC	HP	KS	KC	KP			
1994-1997	Regulated	28	12	0	4	17	14	19	19	18	19		
	Unregulated	59	32	0	24	40	36	33	38	40	25		
1997-2000	Regulated	29	3	3	1	6	2	4	3	3	12		
	Unregulated	50	23	0	16	25	19	13	16	14	22		
2000-2003	Regulated	37	11	0	8	8	9	10	10	10	12		
	Unregulated	50	23	1	13	22	22	15	18	20	16		
2003-2006	Regulated	42	27	0	9	23	22	16	21	19	13		
	Unregulated	54	29	0	19	27	28	22	21	20	21		
2006-2009	Regulated	47	10	0	2	7	8	3	10	8	8		
	Unregulated	85	22	3	14	21	20	13	11	10	18		
Total		481	192	7	110	196	180	148	167	162	166		

Note: Entries are the number of firms with asymmetric returns according to the tests proposed by Lisi (the skewness coefficient \widehat{S} with bootstrapped standard errors), Chen and Lin (the H and K tests), and Maasoumi and Racine (the \widehat{S}_ρ test), all at a 10% level of significance. In this table only the \widehat{S} test is capable of distinguishing positive (right-skewed) from negative (left-skewed) asymmetry. Its results are reported with indicative (+) and (-) signs.

Table 3.4: Results of the asymmetry tests described in Section 3.2.2, at 5% and 10% nominal sizes

Period	Sampled Firms		Conditional Tests at 5%			Conditional Tests at 10%				
	Type	Quantity	$\hat{\xi}_i(+)$	$\hat{\xi}_i(-)$	CS	\hat{S}_ρ	$\hat{\xi}_i(+)$	$\hat{\xi}_i(-)$	CS	\hat{S}_ρ
1994-1997	Regulated	28	12	0	19	17	17	0	21	17
	Unregulated	59	16	1	40	21	23	1	43	26
	Regulated	29	2	0	10	6	6	0	13	10
1997-2000	Unregulated	50	12	0	35	15	15	1	38	21
	Regulated	37	4	1	10	4	5	1	12	8
2000-2003	Unregulated	50	8	0	26	12	11	0	29	19
	Regulated	42	16	0	15	12	22	0	21	24
2003-2006	Unregulated	54	17	0	28	23	22	0	32	26
	Regulated	47	7	2	12	7	10	3	15	9
2006-2009	Unregulated	85	12	0	36	22	24	2	43	29
	Total	481	106	4	231	139	155	8	267	189

Note: Entries are the number of firms with asymmetric returns according to the tests proposed by Lisi (the skewness coefficient \hat{S} with bootstrapped standard errors), Bai and Ng (the conditional symmetry *CS* test) and the QMLE estimate of the asymmetry parameter of the GARCH residual $\hat{\xi}_i$, at confidence levels 5% and 10%. In this table only the $\hat{\xi}_i$ test is capable of distinguishing positive (right-skewed) from negative (left-skewed) asymmetry, and their results are reported with indicative (+) and (-) signs.

Table 3.5: True contingency table under observability of S

	Asymmetric	Symmetric	Total
Regulated	x	$m - x$	n_1
Unregulated	$k - x$	$n - (k - x)$	n_0
Total	k	$m + n - k$	$n_0 + n_1$

Table 3.6: Observed tabulation based on an estimator \hat{S} of S

	Asymmetric	Symmetric
Reg.	$A \equiv x(1 - \beta) + (n_1 - x)\alpha$	$B \equiv (n_1 - x)(1 - \alpha) + x\beta$
Unreg.	$C \equiv (k - x)(1 - \beta) + (n_0 - k + x)\alpha$	$D \equiv (n_0 - k + x)(1 - \alpha) + k\beta$
Total	$k(1 - \beta) + (n_0 + n_1 - k)\alpha$	$(n_0 + n_1 - k)(1 - \alpha) + k\beta$

Table 3.7: Odds ratio by period and asymmetry test performed. Power = 95%

Test	Period (5% significance level)				Period (10% significance level)					
	94-97	97-00	00-03	03-06	06-09	94-97	97-00	00-03	03-06	06-09
\hat{S}	0.43	0.31•	0.71	0.81	0.12•	0.59	0.22•	0.41•	1.66	0.48
HS	0.23•	NA	1.14	0.37	NA	0.06•	NA	0.71	0.40•	NA
HC	0.84	0.09•	0.38•	1.07	0.52	0.74	0.19•	0.24•	1.18	0.35
HP	0.33•	0.10•	0.46	1.18	0.25•	0.56	NA	0.29•	1.00	0.52
KS	1.27	0.86	0.83	0.74	0.44	1.77	0.15•	0.74	0.92	NA
KC	1.11	NA	0.55	1.21	1.38	1.16	0.00•	0.50	1.81	3.95••*
KP	0.91	NA	0.71	1.32	0.90	0.86	0.00•	0.42•	1.36	3.82
\hat{S}_ρ	1.26	0.48	0.69	0.71	0.76	3.26••*	0.92	1.05	0.58	0.63
$\hat{\xi}_i$	2.18*	0.15•	0.89	1.32	1.47	2.76••*	0.46	0.55	1.83	0.88
CS	1.10	0.21•	0.30•	0.47•	0.44•	1.02	0.20•	0.30•	0.68	0.37•
\hat{S}_ρ	2.97••*	0.54	0.21•	0.48	0.39	2.39••*	0.63	0.31•	1.40	0.31•

Notes: Entries are the odds ratio (3.20) associated with a contingency table (see Table 3.5). This table shows results of each asymmetry test considered in sections 3.2.1 and 3.2.2, adjusted for the effect (i.e., size and power) of the statistical test, as shown in Table 3.6. As the level and power of each test is not known, we tested whether the odds ratio was greater than unity for two different rejection probabilities under the null (5% and 10%). Marks that indicate significance are: • = two-sided alternative ($OR \neq 1$), * = one-sided alternative ($OR > 1$).

Table 3.8: Odds ratio by period and asymmetry test performed. Power = 80%

Test	Period (5% significance level)					Period (10% significance level)				
	94-97	97-00	00-03	03-06	06-09	94-97	97-00	00-03	03-06	06-09
\hat{S}	0.43	0.32•	0.79	0.96	0.09•	0.52	0.15•	0.30•	2.31•*	0.56
HS	0.16•	NA	1.19	0.39	NA	0.10•	NA	0.62	0.37•	NA
HC	0.69	0.07•	0.33•	0.98	0.52	0.51	0.12•	0.21•	1.33	0.26•
HP	0.31•	0.08•	0.42•	1.08	0.19•	0.50	NA	0.30•	0.94	0.52
KS	1.32	0.56	0.82	0.74	0.35	2.34*	0.27	0.83	0.85	NA
KC	1.15	NA	0.49	1.40	1.48	1.30	0.00•	0.53	1.93*	5.53•*
KP	0.86	NA	0.79	1.51	0.71	0.84	0.00•	0.45	1.56	4.88•*
\hat{S}_ρ	1.18	0.53	0.60	0.65	0.79	5.35•*	0.88	1.02	0.65	0.66
ξ_i	2.09*	0.12•	0.64	1.51	1.77	3.13•*	0.34	0.41	1.83	0.82
CS	0.83	0.09•	0.26•	0.40•	0.39•	1.47	0.04•	0.23•	0.56	0.35•
\hat{S}_ρ	4.30•*	0.51	0.25•	0.45•	0.37•	2.56•*	0.62	0.29•	1.59	0.29•

Notes: Entries are the odds ratio (3.20) associated with a contingency table (see Table 3.5). This table shows results of each asymmetry test considered in sections 3.2.1 and 3.2.2, adjusted for the effect (i.e., size and power) of the statistical test, as shown in Table 3.6. As the level and power of each test is not known, we tested whether the odds ratio was greater than unity for two different rejection probabilities under the null (5% and 10%). Marks that indicate significance are: • = two-sided alternative ($OR \neq 1$), * = one-sided alternative ($OR > 1$).

Table 3.9: Odds ratio by period and asymmetry test performed. Power = 65%

Test	Period (5% significance level)					Period (10% significance level)				
	94-97	97-00	00-03	03-06	06-09	94-97	97-00	00-03	03-06	06-09
\hat{S}	0.36 [•]	0.25 [•]	0.67	0.83	0.07 [•]	0.40 [•]	0.15 [•]	0.24 [•]	10.30 ^{••}	0.50
HS	0.19 [•]	NA	1.06	0.34 [•]	NA	0.06 [•]	NA	0.71	0.34 [•]	NA
HC	0.66	0.05 [•]	0.33 [•]	1.01	0.48	NA	0.11 [•]	0.17 [•]	1.63	0.25 [•]
HP	0.21 [•]	0.05 [•]	0.41 [•]	1.03	0.24 [•]	0.19 [•]	NA	0.23 [•]	1.01	0.45
KS	2.00	0.41	0.91	0.70	0.29	NA	0.18 [•]	0.69	0.74	NA
KC	1.43	NA	0.44 [•]	1.30	1.04	NA	0.00 [•]	0.46	2.59 ^{••}	5.40 ^{••}
KP	0.83	NA	0.67	1.51	0.9	NA	0.00 [•]	0.34 [•]	1.93 [*]	3.95 ^{••}
\hat{S}_ρ	1.41	0.45	0.63	0.57	0.57	NA	0.95	1.02	0.57	0.59
ξ_i	2.78 ^{••}	0.08 [•]	0.71	1.51	1.69	10.01 ^{••}	0.40 [•]	0.43	2.35 ^{••}	0.78
CS	NA	NA	0.18 [•]	0.32 [•]	0.31 [•]	NA	NA	0.10 [•]	0.36 [•]	0.25 [•]
\hat{S}_ρ	12.25 ^{••}	0.58	0.29 [•]	0.37 [•]	0.40 [•]	8.14 ^{••}	0.59	0.28 [•]	2.73 ^{••}	0.27 [•]

Notes: Entries are the odds ratio (3.20) associated with a contingency table (see Table 3.5). This table shows results of each asymmetry test considered in sections 3.2.1 and 3.2.2, adjusted for the effect (i.e., size and power) of the statistical test, as shown in Table 3.6. As the level and power of each test is not known, we tested whether the odds ratio was greater than unity for two different rejection probabilities under the null (5% and 10%). Marks that indicate significance are: [•] = two-sided alternative ($OR \neq 1$), ^{*} = one-sided alternative ($OR > 1$).

Table 3.10: Proportion of tables with odds ratio different from unity

Panel A: Unconditional Asymmetry						
	Alternative: $OR \neq 1$			Alternative: $OR > 1$		
	$\beta = 5\%$	$\beta = 20\%$	$\beta = 35\%$	$\beta = 5\%$	$\beta = 20\%$	$\beta = 35\%$
$\alpha = 5\%$	22%	25%	33%	0%	0%	0%
$\alpha = 10\%$	36%	39%	58%	6%	17%	16%
Panel B: Conditional Asymmetry						
	Alternative: $OR \neq 1$			Alternative: $OR > 1$		
	$\beta = 5\%$	$\beta = 20\%$	$\beta = 35\%$	$\beta = 5\%$	$\beta = 20\%$	$\beta = 35\%$
$\alpha = 5\%$	47%	60%	69%	13%	13%	15%
$\alpha = 10\%$	47%	47%	77%	13%	13%	31%

Note: Entries are the proportion of odds ratios from Table ?? that are statistically different from unity, for each pair (α, β) .

Table 3.11: Randomness in the sequence of $\{r_{(i)}\}$ ordered by asymmetry

Test	Period				
	1994-1997	1997-2000	2000-2003	2003-2006	2006-2009
Panel A: Unconditional Tests					
\widehat{S}	0.23	0.53	0.01	0.40	0.85
HS	0.02	0.25	0.89	0.25	0.95
HC	0.07	0.05	0.02	0.56	0.92
HP	0.93	0.08	0.29	0.40	0.68
KS	0.89	0.25	0.05	0.25	0.54
KC	0.84	0.00	0.78	0.92	0.54
KP	0.40	0.02	0.54	0.25	0.80
\widehat{S}_ρ	0.16	0.08	0.29	0.72	0.54
Panel B: Conditional Tests					
$\widehat{\xi}_i$	0.31	0.79	0.11	0.01	0.25
CS	0.84	0.13	0.37	0.64	0.46
\widehat{S}_ρ	0.40	0.01	0.45	0.32	0.46

Note: Entries are the one-sided p -values of the runs test on $\{r_{(i)}\}$ for each asymmetry measure and sub-period, where p -values lower than 0.10 are displayed in boldface. The alternative hypothesis is too few runs.

Table 3.12: Results of the Z_A test for equality of empirical distributions.

Test	Period				
	1994-1997	1997-2000	2000-2003	2003-2006	2006-2009
Panel A: Unconditional Tests					
\widehat{S}	0.01	0.00	0.12	0.68	0.22
HS	0.10	0.02	0.67	0.07	0.02
HC	0.32	0.00	0.13	0.00	0.05
HP	0.16	0.00	0.64	0.02	0.54
KS	0.83	0.44	0.61	0.77	0.77
KC	1.00	0.00	0.83	0.22	0.47
KP	0.27	0.00	0.78	0.22	0.76
\widehat{S}_ρ	0.94	0.00	0.01	0.64	0.53
Panel B: Conditional Tests					
$\widehat{\xi}_i$	0.06	0.03	0.05	0.31	0.10
CS	0.53	0.00	0.05	0.25	0.55
\widehat{S}_ρ	0.05	0.00	0.14	0.17	0.17

Note: Entries are the two-sided p -values of the Z_A test as defined in (3.25c) for the hypothesis (3.2), where p -values lower than 0.10 are displayed in boldface.

Table 3.13: Results of the \widehat{D} test for equality of density estimates

Test	Period				
	1994-1997	1997-2000	2000-2003	2003-2006	2006-2009
Panel A: Unconditional Tests					
\widehat{S}	0.01	0.01	0.06	0.41	0.57
HS	0.06	0.03	0.43	0.01	0.19
HC	0.15	0.04	0.06	0.01	0.56
HP	0.22	0.00	0.06	0.00	0.10
KS	0.40	0.05	0.50	0.43	0.15
KC	0.68	0.01	0.24	0.21	0.07
KP	0.04	0.00	0.51	0.20	0.07
\widehat{S}_ρ	0.12	0.02	0.03	0.13	0.18
Panel B: Conditional Tests					
$\widehat{\xi}_i$	0.07	0.01	0.03	0.16	0.08
CS	0.41	0.00	0.17	0.09	0.33
\widehat{S}_ρ	0.01	0.00	0.02	0.04	0.01

Note: Entries are the two-sided p -values of the \widehat{D} test as defined in (3.26) for the hypothesis (3.2), where p -values lower than 0.10 are displayed in boldface.

Table 3.14: Results of the Wilcoxon test for location shift

Test	Period				
	1994-1997	1997-2000	2000-2003	2003-2006	2006-2009
Panel A: Unconditional Tests					
$ \widehat{S} $	0.03	0.96	0.83	0.33	0.79
HS	0.91	0.97	0.90	0.65	0.91
HC	0.80	1.00	0.92	0.59	0.72
HP	0.92	1.00	0.90	0.69	0.82
KS	0.27	0.81	0.87	0.38	0.76
KC	0.50	1.00	0.85	0.29	0.08
KP	0.57	1.00	0.88	0.35	0.24
\widehat{S}_ρ	0.74	1.00	0.99	0.56	0.86
Panel B: Conditional Tests					
$ \widehat{\xi}_i $	0.09	0.92	0.77	0.06	0.95
CS	0.56	1.00	0.99	0.92	0.91
\widehat{S}_ρ	0.92	1.00	0.98	0.75	0.96

Note: Entries are one-sided p -values of the Wilcoxon rank-sum test for the hypothesis (3.4), where p -values lower than 0.10 are displayed in boldface.

Table 3.15: Results of the tests for scale shift

Test	Period				
	1994-1997	1997-2000	2000-2003	2003-2006	2006-2009
Panel A: Unconditional Symmetry					
Ansari-Bradley	0.00	0.08	0.16	0.77	0.79
Mood	0.01	0.11	0.09	0.76	0.91
Panel B: Conditional Symmetry					
Ansari-Bradley	0.85	0.45	0.17	0.91	0.16
Mood	0.94	0.55	0.06	0.93	0.19

Note: Entries are the one-sided p -values of the Ansari-Bradley and Mood tests for the hypothesis (3.3). The unconditional asymmetry metric used is the skewness statistic $\widehat{S}(\eta_t)$ as proposed by Lisi, and conditional asymmetry is measured by the parameter $\widehat{\xi}_i$ of the GARCH innovation of returns. The results for the 1997-2000 period must be interpreted with caution, as $f_{S|R}(s|r = 0)$ and $f_{S|R}(s|r = 1)$ are likely to have a different location parameters during this period as suggested by the results in Table 3.14. The distribution theory of the scale tests used is known only when the location of the two populations is the same or is known in advance and removed from the test samples.

Appendix A

Proof for Proposition 2.3.1

I present a backward induction proof for the SPNE that we are interested in. The approach adopted rules out profitable deviations from the strategies considered above.

1) *Second Period*

First, if there is no innovation then only consumers of type L are remaining in the market in period two, since type H have already bought the good (independently of the quality level) during $t = 1$ and they will not purchase the same product again. It cannot be the case that only type L consumers buy the product in period one, whereas type H forgoes more utility from postponing purchase than type L . For this reason, given the strategy played by the monopolist in period $t = 1$, in the second period she/he sets $q_2 = \bar{q}$ to attract type H back to the market. A necessary condition to guarantee that type H will be willing to buy the second generation is $\bar{v}^H - p_2 > \underline{v}^H$. Basically, type H is comparing the utility gain of switching from the first generation product to the next generation during $t = 2$.

We can guarantee that $q_2 = \underline{q}$ is not a profitable deviation if there exists p_2 and p'_2 such that:

$$\pi_2 = p_2(1 + \mu) - \frac{\gamma(\bar{q} - \underline{q})^2}{1 + \mu\beta_H} \geq p'_2 = \pi'_2 \quad (\text{A.1})$$

The left hand side of (A.1) is the profit that the monopolist gets in the second period in case that he sells the product to both types of consumer, while

the right hand side is the profit that he makes in $t = 2$ in the case that only type L is in the market. Furthermore, insofar as $p'_2 \leq \underline{v}^L$ and $p_2 \leq \bar{v}^L$, a sufficient condition is given by

$$\bar{v}^L \geq p_2 \geq \frac{\underline{v}^L}{(1 + \mu)} + \frac{\gamma(\bar{q} - \underline{q})^2}{(1 + \mu)(1 + \mu\beta_H)} > 0 \quad (\text{A.2})$$

Under the original set of strategies, type L consumer surplus is equal to $\bar{v}^L - p_2$, and under the monopolist deviation strategy it is equal to the $\max\{\underline{v}^L - p'_2, 2\underline{v}^L - p_1\}$. But the monopolist will extract all type's L surplus if they are the only type left, $p'_2 = \underline{v}^L$, while setting $p_1 = \underline{v}^H$. Together, those restrictions imply that $\max\{\underline{v}^L - p'_2, 2\underline{v}^L - p_1\} = \max\{0, 2\underline{v}^L - \underline{v}^H\}$.

Let's further assume that $\underline{v}^L < \underline{v}^H/2$. This inequality says that the benefit type H gets from owning the product in the first period is more than twice the benefit that type L would get. This assumption is not intuitively challenging. If low type cannot properly deal with product's flaws they will rationally derive a low utility level from the \underline{q} version of the durable good. In this case, type L consumers foresee what the monopolist's movement will be and they will buy the product in the second period. As a result, the monopolist deviation $q_2 = \underline{q}$ will not be profitable if $\pi'_2 < 0$, i.e., if $\underline{v}^L < \gamma\underline{q}^2 + \phi\mu/\underline{q}$ (instead of $\pi_2 > 0$, provided that (A.2) holds).¹

2) First Period

Type L will not deviate if $2\underline{v}^L - p_1 \leq \bar{v}^L - p_2 \Rightarrow p_2 \leq \bar{v}^L - 2\underline{v}^L + p_1$, while type H will not postpone his purchase if $\underline{v}^H + \bar{v}^H - (p_1 + p_2) \geq \bar{v}^H - p_2 \Rightarrow \underline{v}^H - p_1 \geq 0$ (a condition that is not trivially satisfied, since we are dealing with a durable good and consequently p_1 can go up to $2\underline{v}^H$). In addition, type H will not stick to the previous generation as long as $\underline{v}^H + \bar{v}^H - (p_1 + p_2) \geq 2\underline{v}^H - p_1 \Rightarrow \bar{v}^H - p_2 \geq \underline{v}^H$ (one of the necessary conditions that prevents the monopolist from deviating and choosing $q_2 = \underline{q}$). Since the monopolist is setting $p_1 = \underline{v}^H$ and fixing $q_2 = \bar{q}$, type H has no incentive to deviate.

¹In order to simplify the analysis, without loss of generality I consider that the damage cost is incurred during the first period.

The possible deviation for the monopolist in period 1 is to set $q_1 = \bar{q}$ and then to choose whether to sell to both types in $t = 1$ or discriminate price intertemporally to skim the demand selling a homogeneous good of quality \bar{q} . However, our assumption above stating that market power matters precludes the former possibility from happening, and hence the monopolist's deviation will not be profitable if there exists p_2 such that

$$\underline{v}^H \mu + p_2(1 + \mu) - \gamma \left(\underline{q}^2 + \frac{(\bar{q} - \underline{q})^2}{1 + \mu\beta_H} \right) - \frac{\phi\mu}{\underline{q}} \geq 2\bar{v}^H \mu + \bar{v}^L - \gamma\bar{q}^2 - \frac{\phi\mu}{\bar{q}}$$

If that is the case, price and quality discrimination will happen simultaneously if

$$p_2 \geq [(2\bar{v}^H - \underline{v}^H)\mu + \bar{v}^L - \underbrace{\gamma \left(\frac{2(\bar{q}\underline{q} - \underline{q}^2) + \mu\beta_H(\bar{q}^2 - \underline{q}^2)}{1 + \mu\beta_H} \right)}_{\text{(I)}} + \underbrace{\frac{\phi\mu(\bar{q} - \underline{q})}{\bar{q}\underline{q}}}_{\text{(II)}}] / (1 + \mu)$$

Since $\bar{q}\underline{q} \geq \underline{q}^2$ the term **(I)** is positive. Its subtraction represents the QA/R&D cost amortization promoted by the consumers' feedback. The term **(II)** is also positive. It enters additively and represents the reputation damage inflicted to the firm. As long as the above conditions are verified that is no incentive for any player to deviate.

Appendix B

Moments of the Return Process

When y_t is written as a function of current and past shocks,

$$y_t = f(\epsilon_t, \epsilon_{t-1}, \dots),$$

the representation in (3.5) can be formalized as a first-order Taylor expansion of $f(\cdot)$ around $\epsilon_t = 0$ given past shocks $\{\epsilon_{t-1}, \epsilon_{t-2}, \dots\}$, and yields

$$y_t = f(0, \epsilon_{t-1}, \epsilon_{t-2}, \dots) + f_1(0, \epsilon_{t-1}, \epsilon_{t-2}, \dots)\epsilon_t, \quad (\text{B.1})$$

where $f_1(\cdot)$ is the derivative of f with regards to ϵ_t . From (B.1) it is easy to derive the conditional moments of y_t . The conditional mean is given by

$$\mu_t \equiv \mathbb{E}_{t-1}y_t = f(0, \epsilon_{t-1}, \epsilon_{t-2}, \dots),$$

and the conditional variance is

$$h_t^2 \equiv \mathbb{E}_{t-1}(y_t - \mathbb{E}_{t-1}y_t)^2 = f_1(\cdot)^2 \mathbb{E}_{t-1}\epsilon_t^2 = f_1(\cdot)^2.$$

The specification in (B.1) also implies that all higher-order conditional moments of y_t are linked to the function $f_1(\cdot)$ and to the error process $\{\epsilon_t\}$. For all $k \geq 2$,

$$\mathbb{E}_{t-1}(y_t - \mathbb{E}_{t-1}y_t)^k = f_1(\cdot)^k \mathbb{E}_{t-1}\epsilon_t^k = h_t^k \mathbb{E}\epsilon_t^k,$$

where in the last equality the conditional expectation was substituted by the unconditional one as ϵ_t is a i.i.d. process by assumption. We could go one step further and form a second-order expansion of $f(\cdot)$ around $\epsilon_t = 0$ by adding the

term $(1/2)f_{11}(0, \epsilon_{t-1}, \epsilon_{t-2}, \dots)\epsilon_t^2$ to (B.1), where $f_{11}(\cdot)$ is the second derivative of f w.r.t. ϵ_t . The potential gain in flexibility however has a high price in terms of tractability, as the conditional variance and all higher-order moments of $\{y_t\}$ will depend on cross-products of powers of two unknown functions, $h(\cdot)$ and $f_{11}(\cdot)$, and such an extension does not seem very promising¹.

The above first-order approximation implies that the conditional skewness of $\{y_t\}$ depends solely on the *unconditional* third moment of the disturbance,

$$\frac{\mathbb{E}_{t-1}(y_t - \mathbb{E}_{t-1}[y_t])^3}{[\mathbb{E}_{t-1}(y_t - \mathbb{E}_{t-1}[y_t])^2]^{3/2}} = \frac{h_t^3 \mathbb{E}_{t-1} \epsilon_t^3}{[h_t^2 \mathbb{E}_{t-1} \epsilon_t^2]^{3/2}} = \mathbb{E} \epsilon_t^3. \quad (\text{B.2})$$

In (B.2) it becomes clear that symmetry of ϵ_t is not an innocuous assumption, as it implies conditional symmetry of the return series y_t .

The unconditional moments of y_t in model in (3.5) are not as simple as the conditional ones. The unconditional mean is

$$\mathbb{E}y_t = \mathbb{E}\mu_t,$$

the unconditional variance is

$$\begin{aligned} \mathbb{E}[(y_t - \mathbb{E}y_t)^2] &= \mathbb{E}[y_t^2 - 2y_t\mathbb{E}\mu_t + (\mathbb{E}\mu_t)^2] \\ &= \mathbb{E}[(\mu_t + \eta_t)^2 - 2(\mu_t + \eta_t)\mathbb{E}\mu_t + (\mathbb{E}\mu_t)^2] \\ &= \mathbb{E}[\mu_t^2 + \eta_t^2 - 2\mu_t\mathbb{E}\mu_t + (\mathbb{E}\mu_t)^2] \\ &= \mathbb{E}\mu_t^2 - (\mathbb{E}\mu_t)^2 + \mathbb{E}\eta_t^2 \\ &= \text{Var } \mu_t + \mathbb{E}\eta_t^2 = \text{Var } \mu_t + \mathbb{E}h_t^2 \quad \blacksquare \end{aligned}$$

and the unconditional third moment is

$$\begin{aligned} \mathbb{E}[(y_t - \mathbb{E}y_t)^3] &= \mathbb{E}[y_t^3 - 3y_t^2\mathbb{E}\mu_t + 3y_t(\mathbb{E}\mu_t)^2 - (\mathbb{E}\mu_t)^3] \\ &= \mathbb{E}[(\mu_t + \eta_t)^3 - 3(\mu_t + \eta_t)^2\mathbb{E}\mu_t + 3(\mu_t + \eta_t)(\mathbb{E}\mu_t)^2 - (\mathbb{E}\mu_t)^3] \\ &= \mathbb{E}[\mu_t^3 - 3\mu_t^2\mathbb{E}\mu_t + 3\mu_t(\mathbb{E}\mu_t)^2 - (\mathbb{E}\mu_t)^3] + 3\mathbb{E}[\mu_t\eta_t^2 - \mathbb{E}\mu_t\eta_t^2] \\ &\quad + \mathbb{E}[3(\mathbb{E}\mu_t)^2\eta_t + 3\mu_t^2\eta_t - 6\mu_t\eta_t\mathbb{E}\mu_t] + \mathbb{E}\eta_t^3 \\ &= \mathbb{E}[(\mu_t - \mathbb{E}\mu_t)^3] + 3\text{Cov}(\mu_t, \eta_t^2) + \mathbb{E}\eta_t^3 \quad \blacksquare \end{aligned}$$

¹Identification in such a model would quickly become a nightmare. For instance, the conditional variance would be given by $f_1(\cdot)^2 + 2f_1(\cdot)f_{11}(\cdot) + f_{11}(\cdot)^2\mathbb{E}_{t-1}\epsilon_t^4$, and the conditional third moment would be $f_1(\cdot)^3\mathbb{E}_{t-1}\epsilon_t^3 + 3f_1(\cdot)^2f_{11}(\cdot)\mathbb{E}_{t-1}\epsilon_t^4 + 3f_1(\cdot)f_{11}(\cdot)^2\mathbb{E}_{t-1}\epsilon_t^5 + f_{11}(\cdot)^3\mathbb{E}_{t-1}\epsilon_t^6$.

where $\mathbb{E}[3(\mathbb{E}\mu_t)^2\eta_t + 3\mu_t^2\eta_t - 6\mu_t\eta_t\mathbb{E}\mu_t] = 0$ since $\eta_t = h_t\epsilon_t$ and ϵ_t is an i.i.d. shock with unit variance. The unconditional skewness of y_t comes from three sources: the conditional mean, which may be skewed by itself or linearly related to the squared disturbance; the conditional standard deviation, which may or may not respond asymmetrically to shocks; and from disturbance itself, which may be asymmetrically distributed. For this reason, to investigate asymmetry in the distribution of returns, one must be careful when imposing restrictions on the three major ingredients of the model (the conditional mean and variance and the error process), since they have a direct influence on the hypothesis being tested. In particular, the variance equation and the disturbance process should be as general as possible, which is the reason we choose the specifications in (3.12) and (3.13).

Appendix C

The Skewed-t Distribution

Fernandez and Steel (1998) show that for any density $f(\cdot)$ continuous, symmetric at zero and unimodal, the transformation

$$g(z | \xi) = \frac{2}{\xi + 1/\xi} [f(z/\xi)\mathbb{I}_{\{z \geq 0\}} + f(z\xi)\mathbb{I}_{\{z < 0\}}] \quad (\text{C.1})$$

for any $\xi \in (0, \infty)$ is a density based on $f(\cdot)$ that retains unimodality at zero for all ξ but loses symmetry whenever $\xi \neq 1$. We take f to be the standardized Student-t density,

$$f(z_t) = \frac{\Gamma((\nu + 1)/2)}{\Gamma(\nu/2)\sqrt{(\nu - 2)\pi}} \left(1 + \frac{z_t^2}{\nu - 2}\right)^{-(\nu+1)/2}. \quad (\text{C.2})$$

The symmetry parameter ξ is proportional to the relative mass of positive values, as $P(x \geq 0 | \xi)/P(x < 0 | \xi) = \xi^2$. The sign of $\log(\xi)$ indicates the sign (or direction) of the asymmetry in $g(\cdot)$: when $\xi \in (0, 1)$, $\log(\xi) < 0$ and $g(\cdot)$ is negatively (of left) skewed; and when $\xi \in (1, \infty)$, $\log(\xi) > 0$ and $g(\cdot)$ is positively (of right) skewed. Figure C.1 illustrates the role of ξ in controlling the skewness of $g(\cdot)$.

The (uncentered) moments of z under $g(\cdot)$ are given by

$$\mathbb{E}[z^r | \xi] = \int_0^\infty s^r 2f(s) ds \frac{\xi^{r+1} + (-1)^r \xi^{-(r+1)}}{\xi + 1/\xi}. \quad (\text{C.3})$$

The skewness coefficient (3.6) is a complicated function under the previous formula, but the asymmetry measure introduced by Arnold and Groenenveld (1995), one




Figure C.1: Skewed versions of a Student-t Density

minus two times the probability mass left of the mode, can be calculated as

$$AG(x|\xi) = \frac{\xi^2 - 1}{\xi^2 + 1}, \quad (\text{C.4})$$

and displays several desirable properties: it is an increasing function of ξ ; preserves the convex orderings of distributions of van Zwet (1964) when $f(\cdot)$ is differentiable; takes values on $(-1, 1)$ and vanishes when $\xi = 1$. Therefore, (C.4) is a bounded measure of asymmetry, with a very intuitive interpretation.

Bibliography

- [1] ADAMS, A., DI BENEDETTO, C. A., AND CHANDRAN, R. Can you reduce your package size without damaging sales? *Long Range Planning* 24, 4 (1991), 86–96.
- [2] ALEXANDER, I., MAYER, C., AND WEEDS, H. Regulatory structure and risk and infrastructure firms: an international comparison. Policy Research Working Paper 1698, The World Bank, 1996.
- [3] ANDERSON, G. The power and size of nonparametric tests for common distributional characteristics. *Econometric reviews* 20, 1 (2001), 1–30.
- [4] ANGRIST, J., AND PISCHKE, J.-S. The Credibility Revolution in Empirical Economics: How Better Research Design is Taking the Con out of Econometrics. *Journal of Economic Perspectives* 24, 2 (2010), 3–30.
- [5] BAI, J., AND NG, S. A consistent test for conditional symmetry in time series models. *Journal of Econometrics* 103, 1-2 (2001), 225–258.
- [6] BAI, J., AND NG, S. Tests for skewness, kurtosis and normality for time series data. *Journal of Business and Economic Statistics* 23, 1 (2005), 49–60.
- [7] BARONE-ADESI, G. Arbitrage equilibrium with skewed asset returns. *The Journal of Financial and Quantitative Analysis* 20, 3 (1985), 299–313.
- [8] BERA, A., AND PREMARATNE, G. Adjusting the tests for skewness and kurtosis for distributional misspecifications. Working Paper 01-0116, Dept. of Economics, University of Illinois, 2001.
- [9] BHATTACHARYYA, M., MISRA, N., AND KODASE, B. MaxVaR for Non-normal and Heteroskedastic Returns. *Quantitative Finance Forthcoming* (2009).
- [10] BOWLEY, A. L. *Elements of Statistics*, 6th ed. New York: Scribner, 1937.
- [11] BRANNAS, K., AND DE GOOIJER, J. Asymmetries in conditional mean and variance: modelling stock returns by asMA-asQGARCH. *Journal of Forecasting* 23, 3 (2004), 155–171.

- [12] BULOW, J. Durable-goods monopolists. *The Journal of Political Economy* (1982), 314–332.
- [13] CAMPBELL, J., LO, A., AND MACKINLAY, A. *The Econometrics of Financial Markets*. Princeton University Press, Princeton, New Jersey, 1996.
- [14] CHALFANT, J. A. A Globally Flexible, Almost Ideal Demand System. *Journal of Business & Economic Statistics* 5, 2 (1987), pp. 233–242.
- [15] CHEN, Y. T., AND LIN, C. C. On the robustness of symmetry tests for stock returns. *Studies in Nonlinear Dynamics & Econometrics* 12, 2 (2008), 1591–1631.
- [16] COASE, R. Durability and monopoly. *Journal of Law and Economics* 15 (1972), 143–149.
- [17] CRAMÉR, H. *Mathematical Methods of Statistics*. Princeton University Press, Princeton, 1946.
- [18] DAUGHETY, A. F., AND REINGANUM, J. F. Communicating quality: A unified model of disclosure and signaling. *The RAND Journal of Economics* (2008).
- [19] DAVIDSON, W. N., RANGAN, N., AND ROSENSTEIN, S. Regulation and systematic risk in the electric utility industry: A test of the buffering hypothesis. *The Financial Review* 32, 1 (February 1997), 163–84.
- [20] DE LIMA, P. J. F. On the robustness of nonlinearity tests to moment condition failure. *Journal of Econometrics* 76, 1-2 (1997), 251–280.
- [21] DEATON, A., AND MUELLBAUER, J. An Almost Ideal Demand System. *American Economic Review* 70, 3 (1980), 312–326.
- [22] DING, Z., GRANGER, C. W. J., AND ENGLE, R. F. A long memory property of stock market returns and a new model. *Journal of Empirical Finance* 1, 1 (1993), 83–106.
- [23] ENGLE, R. F., LILIEN, D. M., AND ROBINS, R. P. Estimating time varying risk premia in the term structure: The ARCH-M model. *Econometrica* 55, 2 (1987), 391–407.
- [24] FERNANDEZ, C., AND STEEL, M. F. J. On bayesian modeling of fat tails and skewness. *Journal of the American Statistical Association* 93, 441 (1998).
- [25] FISHER, R. A. On the interpretation of χ^2 from contingency tables, and the calculation of P. *Journal of the Royal Statistical Society* 85, 1 (1922), 87–94.

- [26] GARVER, M. Listening to customers. *American Journal of Business* 16, 2 (2001), 41–54.
- [27] GIBBONS, J. D., AND CHAKRABORTI, S. *Nonparametric Statistical Inference*. Marcel Dekker, 2003.
- [28] GLOSTEN, L. R., JAGANNATHAN, R., AND RUNKLE, D. E. On the relation between the expected value and the volatility of the nominal excess return on stocks. *Journal of Finance* 48, 5 (1993), 1779–1801.
- [29] GOURVILLE, J. T., AND KOEHLER, J. J. Downsizing price increases: A greater sensitivity to price than quantity in consumer markets. 2004.
- [30] GRANGER, C., AND BILLSON, A. Consumers' attitudes toward package size and price. *Journal of Marketing Research* 9, 3 (1972), 239–248.
- [31] GRANGER, C. W., MAASOUMI, E., AND RACINE, J. A dependence metric for possibly nonlinear processes. *Journal of Time Series Analysis* 25, 5 (2004), 649–669.
- [32] GREEN, R., AND ALSTON, J. M. Elasticities in AIDS Models. *American Journal of Agricultural Economics* 72, 2 (1990), pp. 442–445.
- [33] GRIGOLETTO, M., AND LISI, F. Looking for skewness in financial time series. *The Econometrics Journal* 12, 2 (2009), 310 – 323.
- [34] GROUT, P. A., AND ZALEWSKA, A. The impact of regulation on market risk. *Journal of Financial Economics* 80, 1 (2006), 149 – 184.
- [35] GUL, F., SONNENSCHNEIN, H., AND WILSON, R. Foundations of dynamic monopoly and the coase conjecture. *Journal of Economic Theory* 39, 1 (June 1986), 155–190.
- [36] GUPTA, O. K., TANDON, S., DEBNATH, S., AND ROMINGER, A. S. Package downsizing: is it ethical? *Ai Society* 21, 3 (2006), 239–250.
- [37] HANNAN, E. J. The asymptotic theory of linear time-series models. *Journal of Applied Probability* 10, 1 (1973), 130–145.
- [38] HAUSMAN, J. A. *Valuation of new goods under perfect and imperfect competition*. University of Chicago Press, 1997, pp. 209–237.
- [39] HAUSMAN, J. A., AND LEONARD, G. K. The Competitive Effects of a New Product Introduction: A Case Study. *The Journal of Industrial Economics* 50, 3 (2002), pp. 237–263.
- [40] HEINRICH, J. A guide to the Pearson Type IV distribution. Tech. rep., University of Pennsylvania, 2004.

- [41] HINKLEY, D. V. On power transformations to symmetry. *Biometrika* 62, 1 (1975), 101–111.
- [42] HOLGERSSON, H. Robust testing for skewness. *Communications in Statistics-Theory and Methods* 36, 3 (2007), 485–498.
- [43] HOSKEN, D., AND REIFFEN, D. Patterns of Retail Price Variation. *Rand Journal of Economics* 35, 1 (2004), 128–146.
- [44] IPPOLITO, P. M., AND MATHIOS, A. D. Information, Advertising and Health Choices: A Study of the Cereal Market. *The RAND Journal of Economics* 21, 3 (1990), 459–480.
- [45] JANSEN, D. W., AND DE VRIES, C. G. On the frequency of large stock returns: putting booms and busts into perspective. *The Review of Economics and Statistics* 73, 1 (1991), 18–24.
- [46] JOANES, D. N., AND GILL, C. A. Comparing measures of sample skewness and kurtosis. *The Statistician* 47, 1 (1998), 183–189.
- [47] JURCZENKO, E., AND MAILLET, B. The three-moment CAPM: Theoretical foundations and an asset pricing model comparison in a unified framework. In *Developments in Forecast Combination and Portfolio Choice*, C. Dunis, A. Timmermann, and J. Moody, Eds. John Wiley & Sons, 2001, ch. 13, pp. 239–273.
- [48] KIM, T., AND WHITE, H. On more robust estimation of skewness and kurtosis. *Finance Research Letters* 1, 1 (2004), 56–73.
- [49] KOLBE, A. L., TYE, W. B., AND MYERS, S. *Regulatory Risk: Economic Principles and Applications to Natural Gas Pipelines and Other Industries*. Springer, 1993.
- [50] LAPIDUS, R., AND SCHIBROWSKY, J. Aggregate Complaint Analysis. *Journal of Services Marketing* 8, 4 (1994), 50–60.
- [51] LEE, I., AND LEE, J. A Theory of Economic Obsolescence. *Journal of Industrial Economics* 46, 3 (1998), 383–401.
- [52] LI, Q., MAASOUMI, E., AND RACINE, J. S. A nonparametric test for equality of distributions with mixed categorical and continuous data. *Journal of Econometrics* 148, 2 (2009), 186 – 200.
- [53] LI, Q., AND RACINE, J. Nonparametric estimation of distributions with categorical and continuous data. *Journal of Multivariate Analysis* 86, 2 (2003), 266 – 292.

- [54] LI, Q., AND RACINE, J. S. *Nonparametric Econometrics: Theory and Practice*. Princeton University Press, Princeton, NJ, 2007.
- [55] LISI, F. Testing asymmetry in financial time series. *Quantitative Finance* 7, 6 (2007), 687–696.
- [56] LITTLE, J. D. C. Integrated measures of sales, merchandising, and distribution. *International Journal of Research in Marketing* 15, 5 (1998), 473–485.
- [57] LIU, Z. Stackelberg leadership with demand uncertainty. *Managerial and Decision Economics* 26, 5 (2005), 345–350.
- [58] LOMNICKI, Z. A. Test for departure from normality in the case of linear stochastic processes. *Metrika* 4 (1961), 37–62.
- [59] LORETAN, M., AND PHILLIPS, P. C. B. Testing the covariance stationarity of heavy-tailed time series: An overview of the theory with applications to several financial datasets. *Journal of Empirical Finance* 1, 2 (1994), 211 – 248.
- [60] MAASOUMI, E., AND RACINE, J. Entropy and predictability of stock market returns. *Journal of Econometrics* 107, 2 (2002), 291–312.
- [61] MAASOUMI, E., AND RACINE, J. A robust entropy-based test of asymmetry for discrete and continuous processes. *Econometric reviews* 28, 1-3 (2009), 246–261.
- [62] MEGERDICHIAN, A. A. *Identification of Price Effects in Discrete Choice Models of Demand for Differentiated Products*. Ucsd dissertation chapter, 2009.
- [63] MEGERDICHIAN, A. A. *Identification of Price Effects in Demand Systems with an Application to Merger Simulation*. Ucsd dissertation chapter, 2010.
- [64] MOORS, J., WAGEMAKERS, R., COENEN, V., HEUTS, R., AND JANSSENS, M. Characterizing systems of distributions by quantile measures. *Statistica Neerlandica* 50, 3 (1996), 417 – 430.
- [65] NAIR, H. Intertemporal price discrimination with forward-looking consumers: Application to the us market for console video-games. *Quantitative Marketing and Economics* 5, 3 (2007), 239–292.
- [66] NELSON, D. Conditional heteroskedasticity in asset returns: a new approach. *Econometrica* 59, 2 (1991), 347–370.
- [67] NEVO, A. Measuring Market Power in the Ready-to-Eat Cereal Industry. *Econometrica* 69, 2 (2001), 307–342.

- [68] NEVO, A., AND WOLFRAM, C. Why Do Manufacturers Issue Coupons? An Empirical Analysis of Breakfast Cereals. *The RAND Journal of Economics* 33, 2 (2002), 319.
- [69] NEWEY, W. K., AND STEIGERWALD, D. G. Asymptotic bias for quasi-maximum-likelihood estimators in conditional heteroskedasticity models. *Econometrica*: 65, 3 (1997), 587–599.
- [70] NGATCHOU-WANDJI, J. On testing for the nullity of some skewness coefficients. *International Statistical Review* 74, 1 (2006), 47–65.
- [71] NORTON, S. W. Regulation and systematic risk: The case of electric utilities. *Journal of Law and Economics* 28, 3 (1985), 671–686.
- [72] PELTZMAN, S. Toward a more general theory of regulation. *Journal of Law and Economics* 19, 2 (1976), 211–40.
- [73] PÉREZ-ALONSO, A. A bootstrap approach to test the conditional symmetry in time series models. *Computational Statistics & Data Analysis* 51, 7 (2007), 3484–3504.
- [74] PESENDORFER, M. Retail sales: a study of pricing behavior in supermarkets. *The Journal of Business* 75, 1 (2010), 33–66.
- [75] PLYMIRE, J. Complaints as opportunities. *Business Horizons* 34, 2 (1991), 79 –.
- [76] POLITIS, D. N., AND ROMANO, J. P. The stationary bootstrap. *Journal of the American Statistical Association* 89, 428 (1994), 1303–1313.
- [77] PREMARATNE, G., AND BERA, A. Modeling asymmetry and excess kurtosis in stock return data. Working Paper 01-0123, Dept. of Economics, University of Illinois, 2001.
- [78] PREMARATNE, G., AND BERA, A. A test for symmetry with leptokurtic financial data. *Journal of Financial Econometrics* 3, 2 (2005), 169–187.
- [79] QUINN, J. B. Strategic outsourcing: Leveraging knowledge capabilities. (cover story). *Sloan Management Review* 40, 4 (1999), 9 – 21.
- [80] R DEVELOPMENT CORE TEAM. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2009. ISBN 3-900051-07-0.
- [81] RODRIGUEZ, M. J., AND RUIZ, E. GARCH models with leverage effect: differences and similarities. Working paper 09-03, Departamento de Estadística, Universidad Carlos III de Madrid, 2009.

- [82] ROMANO, J. P., AND SIEGEL, A. F. *Counterexamples in Probability and Statistics*. Chapman & Hall, 1986.
- [83] RUBINFELD, D. L. Market Definition with Dierentiated Products: The Post/Nabisco Cereal Merger. *Antitrust Law Journal* 68, 1 (2000), 163–185.
- [84] SCHIBROWSKY, J., AND LAPIDUS, R. Gaining a Competitive Advantage by Analyzing Aggregate Complaints. *Journal of Consumer Marketing* 11, 1 (1994), 15–26.
- [85] SCHMALENSEE, R. Entry Deterrence in the Ready-to-Eat Breakfast Cereal Industry. *The Bell Journal of Economics* 9, 2 (1978), 305.
- [86] SENTANA, E. Quadratic arch models. *The Review of Economic Studies* 62, 4 (1995), 639–661.
- [87] SHEPARD, A. Price Discrimination and Retail Configuration. *Journal of Political Economy* 99, 1 (1991), 30–53.
- [88] SHUM, M. Does Advertising Overcome Brand Loyalty? Evidence from the Breakfast-Cereals Market. *Journal of Economics Management Strategy* 13, 2 (2004), 241–272.
- [89] SINGLETON, J. C., AND WINGENDER, J. Skewness persistence in common stock returns. *Journal of Financial and Quantitative Analisis* 21, 3 (1986), 335–341.
- [90] SKOGLUND, J. *Essays on Random Effects Models and GARCH*. PhD thesis, Stockholm School of Economics, Stockholm, Sweden., 2001.
- [91] STIGLER, G. J. The theory of economic regulation. *Bell Journal of Economics* 2, 1 (1971), 3–21.
- [92] STOKEY, N. Intertemporal Price Discrimination. *The Quarterly Journal of Economics* 93, 3 (1979), 355–371.
- [93] STOKEY, N. Rational Expectations and Durable Goods Pricing. *The Bell Journal of Economics* 12, 1 (1981), 112–128.
- [94] STUART, A., AND ORD, J. *Kendall’s Advanced Theory of Statistics*. Oxford University Press, 1987.
- [95] TERÄSVIRTA, T. An Introduction to Univariate GARCH Models. In *Handbook of Financial Time Series*, T. G. Andersen, R. A. Davis, J.-P. Kreiss, and T. Mikosch, Eds. Springer Berlin Heidelberg, 2009.
- [96] TIROLE, J. *The theory of industrial organization*. MIT press, 1988.

- [97] TSAY, R. S. *Analysis of Financial Time Series*. John Wiley and Sons, 2005.
- [98] VAN ZWET, W. R. *Convex Transformations of Random Variables*. Mathematisch Centrum, Amsterdam, 1964.
- [99] VERHOEVEN, P., AND MCALEER, M. Fat tails and asymmetry in financial volatility models. *Mathematics and Computers in Simulation* 64, 3-4 (2004), 351–361.
- [100] VISCUSI, W. K., VERNON, J., AND HARRINGTON, J. E. *Economics of Regulation and Antitrust*, 4th ed. MIT Press, 2005.
- [101] WHITE, H. A Heteroskedasticity-Consistent Covariance Matrix Estimator and a Direct Test for Heteroskedasticity. *Econometrica* 48, 4 (1980), 817–838.
- [102] WIDDER, D. V. *The Laplace Transform*. Princeton University Press, Princeton, NJ, 1946.
- [103] WUERTZ, D., CHALABI, Y., MIKLOVIC, M., BOUDT, C., AND CHAUSSE, P. *fGarch: Autoregressive Conditional Heteroskedastic Modelling*. Rmetrics, 2009. R package version 2110.80.
- [104] YAO, Q., AND BROCKWELL, P. Gaussian maximum likelihood estimation for arma models i: time series. *Journal of Time Series Analysis* 27, 6 (2006), 857–875.
- [105] ZAKOIAN, J. Threshold heteroskedastic models. *Journal of Economic Dynamics and Control* 18, 5 (1994), 931–955.
- [106] ZHANG, J. Powerful two-sample tests based on the likelihood ratio. *Technometrics* 48, 1 (2006), 95–103.