

Commitment, Vertical Contracts and Dynamic Pricing of Durable Goods

Job Market Paper.

Please find the updated version at www.stanford.edu/~daljord.

Øystein Daljord*

November 3, 2014

Abstract

Resale Price Maintenance is a vertical contract in which a manufacturer sets the retail price. Traditional motivations for RPM are that it can avoid double marginalisation and provide incentives for complementary service provision. This paper explores a new and complementary role for Resale Price Maintenance (RPM) as part of a price skimming strategy. In dynamic markets, RPM can also fix the price path. By fixing the price path, RPM improve upon price skimming strategies by providing commitment to future prices. I analyse the effect of RPM as a commitment device empirically using a legislation change that deemed RPM illegal as a natural experiment. First, using detailed and comprehensive retail sales data from the Norwegian book market, I show that in the absence of RPM, price skimming falters. Prices fall earlier over the lifecycle, and demand shifts from consumers buying early at high prices to later at lower prices. I then turn to quantify the dynamic effects of RPM. A distinction is made between two effects of RPM: it both precludes price competition between retailers, and it coordinates prices over time. The commitment effect is singled out in a series of counterfactual exercises. I first estimate a dynamic demand model. I then evaluate the returns to counterfactual vertical contracts with varying degrees of commitment in a dynamic oligopoly equilibrium model at the estimated parameters. The commitment value of RPM is estimated to 2.5% increase over an uncoordinated industry profit benchmark. The value of horizontal coordination is estimated at 15.3% using the same benchmark.

1 Introduction

This paper shows how Resale Price Maintenance can improve on price skimming strategies. Price skimming can be a powerful price strategy in markets for durable goods. By gradually lowering the prices over time, it lets consumers with high valuations purchase early, while consumers with lower valuations purchase later at lower prices. An obstacle to price skimming arises when forward looking consumers come to expect future discounts. A consumer that expect a future discount will find a current purchase less attractive. Demand shifts towards future lower prices and price skimming falters. The seller is effectively competing with its own future prices. The phenomenon is known as the Coase conjecture (1972) and has received wide attention in the theory literature (Waldman (2003)). Coasian dynamics has recently been studied empirically in a variety of markets, such as college textbooks

*Thanks to my advisor Harikesh Nair, my committee of Lanier Benkard, Tim Bresnahan and Wesley Hartmann. Thanks also to Chris Conlon, Jeremy Fox, Pedro Gardete, Jakub Kastl, James Lattin, Jon Levin, Erik Madsen, Josh Mollner, Sridhar Narayanan, Christian Riis, Stephan Seiler, Ali Yurukoglu and Pavel Zruev for comments. Thanks to Bokdatabasen for access to data, to Henrik Lande and Gjermund Nese with the Norwegian Competition Authority for facilitating access to the sales data. Thanks to Retriever for access to their media archives. Thanks to Mike Saunders and Marcus Edwall for indispensable help with software. Thanks to Shuo Xie for excellent research assistance.

(Chevalier & Goolsbee (2009)), consumer electronics (Conlon (2012)), video games (Nair (2007)), fashion goods (Soysal & Krishnamurthi (2012)) and sports event tickets (Sweeting (2012)).

Stokey (1981) shows that a seller faced with forward looking consumers would prefer to commit to a price path of higher future prices. A consumer that expect higher future prices would find a current purchase more attractive at an otherwise equal current price. An important question for a seller is where commitment power may be found. I first argue that RPM can provide commitment to future prices by being a legally binding contract that regulates the retail price. In particular, industry wide RPM agreements that are commonly found in book markets are shown to have provisions that provide commitment.

I then empirically quantify the returns to RPM as a commitment device in the Norwegian book market where new legislation imposed changes to the vertical contracts. The Norwegian book market is a compelling setting for an empirical study of vertical contracts and dynamic pricing. The Norwegian book industry has an oligopolistic structure of both publishers and retailers. Until the legislation change, the industry had employed a time limited RPM for decades. The publisher fixed the retail price for a limited period, followed by a heavy discount sale, after which retailers were to price books at their discretion. Following an alignment of the Norwegian competition law with its European counterpart in 2004, the industry's use of RPM was deemed illegal. The change in legislation provides an exogenous change in the vertical contracts that help identify its dynamic effects.

Comprehensive retail sales data show two pronounced changes to the sales patterns following the new legislation: prices fell earlier and demand shifted from early at high prices to later at lower prices. The changes in sales patterns both show evidence of forward looking consumers, and that in the absence of RPM, price skimming falters as more demand is served at lower prices. The changes following the new legislation imply that the vertical contracts helped the industry maintain price strategies that retailers were not able to sustain independently. The changes in sales patterns provide the point of departure for the analysis of vertical contracts as part of a price skimming strategy.

I then turn to quantify the commitment effect of RPM. The main empirical challenge lies in that beyond weakening commitment to future prices, RPM also allows retailers to compete on price. Increased retailer competition itself has a material impact on dynamic pricing patterns. To separate the impact of increased retailer competition from the weakening of the commitment, I turn to series of counterfactual exercises. I first estimate a dynamic demand model, and then evaluate various vertical contracts that coordinate pricing across retailers and over time. The counterfactual contracts differ with respect to the level of commitment they offer.

The data are at the title-retailer level and the demand estimation allows for demand substitution between retailers and over time. The discount factor that in part determines the inter temporal substitution is usually assumed in dynamic demand models since the identifying variation is hard to find in non-experimental data (Magnac & Thesmar (2002)). By fixing the price paths, RPM arguably also fixed consumers price expectations. Using the new legislation as an instrument shifting the retailers price strategies, and consequently the consumers price expectations, the discount factor can be estimated off the sales data. A machine learning algorithm is used in the estimation routine to extract information from editorial text reviews and serves to ameliorate the endogeneity induced by otherwise omitted demand side variables.

The estimates of the dynamic demand model enter into a dynamic supply side model of oligopolistic retail competition that allows for various types of vertical

contracts. The focal contract is the time limited RPM that was used in the Norwegian book market. Three reference benchmark counterfactual contracts are then constructed.

- A baseline contract where pricing is neither coordinated over time nor between retailers. The contract is solved as a dynamic oligopoly game where retailers use time consistent price strategies that satisfy Markov Perfect Nash equilibrium conditions.
- Vertical integration without commitment. The vertical unit uses RPM to coordinate prices between retailers, but has no commitment power. The retailers price strategies are jointly time consistent.
- Vertical integration with commitment. The vertical unit uses RPM to both fully coordinate pricing both across retailers and over time.

The first benchmark is intended as a reference against which to measure various levels of price coordination. The third is a contract that exhausts the scope for coordination of prices both across retailers and over time. The difference between the second and the third contract measures the scope commitment beyond retailer coordination. Finally, the time limited contracts employed in the Norwegian market is compared to the counterfactual benchmark. Preliminary results show that the RPM strategy employed in the Norwegian book market comes quite close to the vertically integrating contract with commitment. Rao (2014) does a similar analysis of commitment versus no-commitment price strategies for a related product category: online video content. Using experimental data, she shows returns to commitment strategies on the order of 40%. This both suggests a role for inter temporal price discrimination for relatively small stakes, and also sizeable returns to commitment policies.

This is primarily a paper about vertical restraints and dynamic pricing, yet the findings have implications for an ongoing policy discussion on RPM. RPM has been illegal in many countries and product markets. The last decades there has however been a shift towards increasing acceptance and endorsement of RPM. Though RPM is *per se* illegal in the EU, book markets are often made exemptions. Having been *per se* illegal in the US for about hundred years¹, RPM is now subject to a rule-of-reason.² A rule-of-reason implies that the legal status of RPM depends on its competitive effects.³ From a policy perspective, it is therefore important to understand the competitive effects of RPM in retail markets, yet little is known empirically, in part because RPM has been an illegal contract form. This paper adds to the scarce empirical evidence on RPM. Results in this paper have particular policy relevance by showing that improved price discrimination may be an effect of RPM, an effect that so far has not been documented in the literature. It is for instance not *a priori* clear that more efficient price discrimination is welfare reducing.

The Norwegian market RPM studied in this paper was implemented through a Trade Agreement that included the majority of publishers and retailers. The Trade Agreement is not unique in its kind. Fixed price agreements, either regulated through trade organisations or by government law, are found in other European

¹For an agreement that is *per se* illegal, the existence of the agreement is illegal, regardless of its effects. Under a rule-of-reason, either the plaintiff must show that RPM is anti-competitive, or the defendant must show that RPM is pro-competitive.

²See *Leegin Creative Leather Products, Inc. vs. PSKS, Inc.* 551 US 877. Following *Leegin*, a minimum price restraint was no longer *per se* illegal. A maximum price was absolved from *per se* illegality status 10 years earlier with *State Oil Co. v. Kahn*, 522 US 3. Along with *Continental T.V., Inc vs. GTE Sylvania Inc*, 433 US 36 where vertical restraints on territories on exclusivity clauses were made subject to a rule-of-reason from being *per se* illegal in 1977, these cases collectively represent a significant and consistent change towards an increasingly friendly policy on vertical restraints in the US.

³See Overstreet (1983) for a survey of historical antitrust case studies and price surveys from the Fair Trade era. Ippolito (1991) surveys RPM cases litigated in the immediate years following the Fair Trade repeal in 1975. See Telser (1960) for a contemporary analysis of the Fair Trade institution itself. Following *Leegin*, some evidence has emerged. In the seemingly most comprehensive empirical study, McKay & Smith (2013) use variation in federal precedents across states and a diff-in-diff approach on grocery scanner data before and after *Leegin*. It finds evidence of price increases, but the study does not consider the dynamic pricing of durable goods explicitly.

book markets as well. Other countries that practice fixed price agreements are France, Germany, Austria, The Netherlands, Italy, Spain, Portugal, Greece, Hungary, Israel, Slovenia, Argentina, South Korea, Japan and Mexico, according to the International Publishers Association. The widespread use of RPM make the book markets interesting in their own right.

Fixed price agreements in book markets are typically motivated by a belief that they stimulate bibliodiversity, that is, a culturally diverse literature and a dense distribution network of book stores. Though neither the supply of variety nor the entry and exit of book stores are the focus of this paper, it points to an effect of RPM that can improve industry profitability so as to sustain bibliodiversity. There is however no clear consensus on the use of RPM as a policy tool in book markets. For instance, while RPM is illegal in the UK book market, RPM is government enforced in the French book market.⁴ The variation in policies across countries reflects a divided opinion on the merits of RPM as a policy instrument. Despite the controversy surrounding the use of RPM, the recent liberalisation of the US legislation and the growing European acceptance of RPM renews its relevance as a managerial pricing tool.

The main managerial content of this paper lies in demonstrating how RPM can coordinate downstream price incentives and mitigate Coasian dynamics at plausibly low costs. Commitment can be achieved through other means than RPM, for instance through reputation. For instance, Apple has a fairly consistent policy across of offering very limited discounts across their product line. The consistent policy can build a reputation for not lowering prices that may persuade consumers to purchase early. Acquiring commitment through reputation can however be both costly and hard. JC Penney, a clothing retail chain, famously adopted an Every Day Low Price policy to avoid losing demand to future discounts. The retail chain however failed to convince consumers, and was forced to revert to frequent sales.⁵ To the extent managers can improve on price skimming through writing vertical contracts, it may prove to be a cost efficient strategy.

The vertical restraints analysed in this paper have become popular in digital markets where it is known as the agency pricing model. Under the agency model, the manufacturer sets the retail price and the retailer gets a share of the revenue. The agency pricing model is used in the App Store and on eBay, and it received wide attention following the recent e-books antitrust case against Apple and a set of US publishers.⁶ Coupled with a Most Favoured Nation clause, the publishers along with Apple forced through an industry wide migration from the classic wholesale pricing model to the agency model. The result was an industry pricing model that closely resembles the RPM in the Norwegian market studied in this paper. De los Santos & Wildenbeest (2014) gives an exposition of the antitrust case and show evidence of substantial increases in prices following the adoption of the agency model. According to case documents, the move to the agency model was motivated in part by publisher's concern that Amazon's heavy discount e-book price policy would lead consumers to expect low prices on all types of books. The role of consumers expectations, and how they relate to vertical restraints, is the topic of this paper.

The impact of RPM as a price skimming strategy is quantified by bringing together concepts and frameworks from the so far mostly distinct literatures on RPM and dynamic pricing. There is a rich theory literature on the effects of vertical contracts and channel coordination across the fields of operations research, economics and marketing, see Cachon (2003) for a survey. The literature on channel coordination

⁴Regarding distribution, the UK has experienced sharp decline in the number of bookstores, whereas France has not. See the press release of the UK Booksellers Association of Oct. 3rd 2011, see <http://www.booksellers.org.uk/campaigns/keepbooksonthehighstreet>.

⁵See <http://www.nytimes.com/2013/04/14/business/for-penney-a-tough-lesson-in-shopper-psychology.html>

⁶DOJ Complaint, US vs. Apple Inc., et al. April 11, 2012

in dynamic markets is however relatively slim. Desai et. al. (2004) is closest to the idea in this paper. It considers a dynamic channel coordinating two-part tariff in a two-period durable goods market with forward looking consumers. In contrast, this paper considers a wider set of contract types, and in particular RPM, it allows for an oligopolistic retail market, it specifies the sources of commitment, and it provides empirical evidence on the effects.

There is little empirical evidence to bear on the effects of vertical contracts. Some empirical papers on vertical contracts have emerged the last few years, e.g. Besanko et al (2005) on retail pass-through, Villas-Boas (2007) on identification of unobserved vertical contracts, Asker & Ljungquist (2013) on the impact of vertical integration in investment banking. By adding vertical restraints to the pricing problem, this paper also contributes to the scarce empirical evidence on vertical contracts and supply chain coordination in general. The two empirical papers topically closest to the current paper are Ho et al. (2012) and Mortimer et al. (2008). The first examines full line forcing in the video rental industry, the second studies revenue sharing in the same industry. Neither however consider explicitly the role of vertical contracts for dynamic pricing. Finally, this paper contributes to a small, but growing empirical literature on dynamic oligopoly pricing of durable goods (Conlon (2012), Goettler & Gordon (2011), Gowrisankaran & Rysman (2009)).

The next section describes the price discrimination problem in a dynamic market with an upstream publisher and oligopolistic downstream retailers. The discussion sketches how vertical contracts improve price skimming in markets with forward looking consumers. Sections (3) and (4) describe the vertical contracts employed in the Norwegian book industry, the legislation change and the data. Section (5) shows reduced form evidence of the impact of RPM on prices and sales before and after the legislation change. Section (6) describes the counterfactuals in terms of a demand and a supply side model. The estimation routine is described in sections (8) and (9). The empirical results are reported along with the counterfactual simulations in sections (10) and (11).

2 Price skimming and vertical contracts in oligopolistic markets

This section provides an informal discussion of some of the issues raised in implementation of price skimming strategies in oligopolistic retail markets. It explains how the vertical unit, a publisher and a set of retailers, can use vertical contracts to address these issues. The discussion leads to a description of the RPM contracts used in the Norwegian book industry. An important distinction is made between the effect of RPM on coordinating prices between retailers (horizontal coordination), and the effect on coordinating prices over time (inter temporal coordination). Though conceptually different, both levels of coordination, or lack thereof, affect the vertical unit's ability to price skim.

Some examples are instructive. Suppose first a publisher has a set of books to sell in a market with a fixed number of heterogenous consumers. Each consumer has unit demand. The publisher adopts a price skimming strategy that gradually lowers the price over time. The publisher's dynamic trade off in any period is between lowering the current price and increase the current profits, at the expense of reducing future profits. Lowering the current price may increase the current profits by increasing the current demand, at the expense of tapping into the future demand. The reduced future demand will furthermore be served at lower prices as high valuation consumers are already cleared out of the market. Making the dynamic trade-off, the publisher gradually lowers the price and appropriates the surplus the books generate.⁷

⁷The returns to optimising price skimming strategies can be substantial in dynamic markets. Lazarev (2013) finds that modern airline pricing strategies extract on the order of 90% of consumers surplus.

Adding retailers to the supply chain requires some modifications to the simple price skimming strategy. Suppose the publisher employed a set of competing retailers to sell the books. The dynamic trade-off for any one retailer considering lowering the current price is again between increasing the current profits at the expense of reducing future demand. The future demand is however of less value to the retailer for two reasons. Firstly, the future demand is shared with its rival retailers, which debases its value. Secondly, by lowering the current price, the retailer gets some of the current demand of his rival retailers. In individual pursuit of market shares, the retailers collectively fail to efficiently price skim the market. RPM can eliminate wasteful competition between the retailers by directly controlling the retail prices, where wasteful is understood from the perspective of the vertical unit.⁸

Forward-looking consumers presents the supply side with a further challenge . A forward looking consumer that correctly predicts future discounts will at otherwise equal terms find a current purchase less attractive. Then high valuation demand shifts towards future lower prices and price skimming falters. The phenomenon is known as the Coase conjecture in the durable goods literature, going back to Coase (1972).

To counteract the inter temporal substitution, the publisher could announce that the future will hold no discounts. The announcement may however not be credible. To see that, suppose the consumers believe the announcement. Consumers with a valuation in excess of the price now have no reason to delay the purchase. As soon as the high valuation consumers are cleared out of the market, the seller has an incentive to reduce the price to capitalise on the low valuation demand left in the market. But then the announced future prices are inconsistent with the actual prices. It is hard to imagine that announcing future prices that consistently fail to realise can form part of a long run, viable price skimming strategy. The seller loses profits to competition from its own future pricing, and is left looking for other means to coordinate prices over time.

Suppose now that instead the publisher exposed itself to costly consequences was he not to price along the announced path. The potential consequences can provide the publisher a *commitment* to the announced price path by counteracting the incentive to discount as soon as the high valuation consumers have left the market (Stokey (1981)).

The discussion emphasises three key components to a price skimming strategy in a market with forward looking consumers. Firstly, the seller needs to announce a price path. Secondly, the price path must persuade high valuation consumers to purchase early at high prices. Thirdly, a counteracting incentive is required to ensure the seller prices along the announced price path. In the next section, the RPM agreement in the Norwegian book market is shown to both announce a price path, and have externally enforced sanctions in place to provide the vertical unit commitment to the announced price paths.

3 Policy shock

The book industry in Norway employed a trade agreement ("Agreement") dating back to the 1960s up until the legislation change became effective in May 2005.⁹ The old Agreement was a legally binding contract voluntarily entered between the Association of Booksellers and the Association of Book Retailers that specified the terms of sales in the industry. The restraints of the Agreement had two key compo-

⁸Other restraints can in principle achieve the same outcome. In a companion paper, I develop a vertical contract that sustains the efficient price path through a path of wholesale prices and transfers, see Daljord (2014b). Nair (2007) report that decreasing wholesale price paths are seemingly common in the video game industry.

⁹See "Bokavtalen" in Store Norske Leksikon (Norwegian Encyclopedia) at www.snl.no .

nents. The first was that the publisher specified a time-limited price restraint. The second was an industry coordinated clearance sale following the expiration of the price restraint period. The time limited fixed price and the coordinated clearance sale together trace out a price path.

The Agreement had kept the pricing strategies stable in the industry for decades and was therefore well known to the consumers. The fixed price was often hard printed onto the cover of the book, which served as an announcement of the retail prices to consumers. There was little price promotion in the industry except for the clearance sale, which was trade marked and heavily advertised.

The Agreement also specified arbitration clauses in case of a breach of contract. From Clause 5:

Violation of the provisions of this Agreement may be prosecuted and, if necessary, by any of the two associations, any publisher and any bookstore or combinations of these who through their union are affiliated by the Trade Agreement. Each association further commits to, within the framework of the individual association bylaws, to take appropriate measures against its own members who may be guilty of violations of this Trade Agreement.

Two features of the sanctions stand out. Firstly, the Agreement did not coordinate on a price level, at least not explicitly, but on a price path. Each publisher was at liberty to set any retail price, but once set, the retailers had to respect the fixed price over the restraint period.¹⁰

Secondly, though the agreement regulates a bilateral agreement between a publisher and a retailer on the shape of the price path, the Agreement exposes the vertical unit to threats of legal action by rival firms if the unit was it to deviate from the specified price path. Beyond allowing for legal actions of rival firms, the Agreement also allowed the Associations to meter out further punishments within the confines of each associations bylaws. The arbitration clause provides means of external enforcement of a bilateral agreement, consistent with commitment. The threat of being taken to court, or made subject to other punishments, can counteract the incentive to deviate from the announced price path. We will however see in the empirical section that by and large, the Agreement was respected.

As part of European legislative integration, the Norwegian competition law was aligned with its European Union counterpart in 2004. Following the legislation change, the Norwegian Competition Authority deemed the Agreement unlawful and called for abolishment. European integration is a political process that evolves independently of developments in the Norwegian book industry, and hence the legislation change can be considered exogenous. The industry voiced strong and united opposition against the new legislation, suggesting the Agreement helped solve an industry coordination problem. The Association of Booksellers, the Association of Publishers and the Association of Authors rallied together against the new legislation and called for exemption from the competition law.¹¹ A public debate ensued and a political compromise was reached.¹²

The new legislation resulted in an exogenous change to the vertical restraints in the

¹⁰The Trade Agreement has been suspected of facilitating horizontal collusion among publishers. The idea is that with RPM, it is easier for publishers to detect deviations on observable retail prices than say on unobservable and flexible wholesale prices, see Jullien & Rey (2007) for one treatment of the argument.

¹¹Exemptions from the competition law can be given for industries that make goods considered to be of particular importance to national identity and is widely allowed for cultural goods, see Canoy & van der Ploeg (2005).

¹²To give some context of the media attention devoted to the new legislation, a search on the keywords 'Book Trade Agreement' in Retriever, a comprehensive Scandinavian media archive, over the period of public debate gives about half the search hits that 'Salt Lake City Olympics' gives over a comparable period at the time of the contemporaneous winter olympics. The numbers can give some perspective of the media interest the new legislation spurred in a nation which is above average preoccupied with winter sports.

industry. The main changes under the new Trade Agreement effective the spring of 2005 were twofold.

- A *shortening* of the price restraint period by eight months, from the year of publication plus one year to the year of publication plus four months.
- A *softening* of the fixed price to a price band. Whereas the RPM under the old regime was a floor and a ceiling, retailers were given discretion to discount the fixed price by up to 12.5% under the new regime.

The changes to the price restraints following the legislation are illustrated in Figure (1).




Figure 1: The price restraints followed calendar time. Some title 1 released early in the year and some other title 2 released later in the year would both have their price restraints lifted at the end of the following year under the old regulation. Following the expiration of the price restraint period, the titles went to the clearance sale with discounts on average in the range of 40% to 50%. The end of the clearance sale marks the end of the typical title lifecycle and most titles sell little. After the legislation change, the same titles 1 and 2 would again have their restraints lifted at the same calendar time, but now May 1st the year after publication rather than December 31st. The clearance sale continued to be held in Spring, but was no longer part of the Agreement itself.

The new vertical restraints implied shorter period of commitment future prices, and a weakening of the restraints themselves. We will see in the coming sections that these exogenous changes in the vertical contracts had substantial impact on both retailers price strategies and the demand patterns.

4 Data

The sales data are scanner data collected from the four largest book retail chains in four month periods over the years 2004 to 2007, bookending the legislation change effective in May 2005. The data make up about 50% of total national sales over the period. The data is aggregated over four months, tertiles, and across stores within each chain. Observations are on title level identified by an Electronic Article Number (EAN) and contains data on a little more than 27000 titles. The EAN identifier allows the sales data to be merged with a comprehensive catalogue of title characteristics provided by Bokdatabasen, an industry logistics company. The catalogue contains data on the fixed price, genres, and various other characteristics such as page counts, edition etc and is used by retailers for logistical purposes and ordering. Prices are calculated as revenue divided by quantity sold in each period for each chain. Price policies were mostly uniform within the chains, according to industry representatives. The summary statistics of the scanner data are given in Table (1).

Table 1: Summary statistics matched scanner data

Variable	Mean	Std. Dev.
retail price	181.51	128.14
quantity	27.58	171.48
year	2005.68	1.08
fixedprice	218.19	129.47
<i>N</i>	1127867	

5 Impact of legislation change

Figure (2) illustrates the window of data and the legislation change. The lifecycle of a title is taken to be about three years. After three years in the market, the sales of most titles are exhausted and all price restraints were lifted both before and after the deregulation. There are two cohorts of full life cycles in the data. These are the titles published in 2004 and the titles published in 2005. I would ideally compare the price paths and demand for a typical title over the lifecycle before and after the legislation change. The data however cover only one year of sales before the legislation.




Figure 2: Data and regulation time line

To establish the sales patterns before and after the legislation change, I assume that the price and sales patterns are comparable across cohorts within a given year. The assumption implies that though individual titles change across years, the mean prices and aggregate sales are drawn each year drawn from the same distribution. I then construct a lifecycle price path before the legislation change by taking the mean price path of new titles in 2004, splice it with the price path of one year old titles in 2004 and lastly, with the two year old titles in 2004.¹³ The resulting price path serves as a measure of the representative price path before the legislation change. Holding the release schedule fixed, we can then make meaningful comparisons of the differences in sales before and after the legislation change.

There is no clear evidence of changes in the release dates of titles after the new legislation. That is somewhat surprising since changes in the duration of the restraint period might change the optimal timing of releases. For instance, publishers might prefer to release titles earlier in the year to benefit from the restraints for longer. Einav (2007) for instance finds evidence of strategic timing of releases in the U.S. motion picture industry. The lack of changes in release dates may be related to the seasonality of demand. For instance, fall is the season for premium fiction. It could be that it is more important to release a title into the market at peak demand than it is to receive a longer lasting protection against future price discounts.

The data comes aggregated in four month periods (tertiles) which are labelled spring, summer and fall. The mean representative prices before and after the legislation change are graphed in Figure (3). The prices are normalised to the fixed price publishers set at the time of publication. A price of 1 implies that a title on average retailed at the fixed price, whereas a price of say 0.5 means a title retailed at 50% discount. The normalisation allows comparison of price paths across different price points. The prices are plotted against time at the tertile periodisation of the the data and contains a total of nine points. Confidence intervals of the means are linearly interpolated between the data points to display the variance. Since standard errors are relatively tight, the confidence intervals are reported at non-conventional levels to display visually discernible variation over time.

The price restraints before and after the legislation change are denoted below

¹³Note that as titles are released over the year of publication, the set of titles in the sample is growing.

the graphs. After, the retailers were allowed to discount the fixed price by up to 12.5% at discretion. The resulting price band is illustrated by the shaded area. The retailers are seen to have largely respected the fixed price policy under the old Agreement. Titles were retailing close to the fixed price in the price restraint period. Towards the end of the restraint period, there are some signs of retailers allowing discounts on the fixed price, on average about 5%. The deviations show that RPM was effectively a minimum price restraint. The restraint period was followed by the industry coordinated clearance sale which saw average discounts on the order of 45%. The clearance sale marks the end of the lifecycle for most titles. The average prices are seen to bounce back somewhat after the clearance sale. The bounce back may have a variety of causes. Firstly, there is a selection of titles that still sell after the sale, and these title may command a higher prices. There may also be seasonal patterns, fall is the premium season. Finally, after the clearance sale, retailers may put the few remaining copies back in the shelf at higher prices to retain variety and the occasional consumer with a higher willingness to pay, somewhat in the spirit of Sobel (1984).




Figure 3: Mean of retail prices normalised to the fixed price.




Figure 4: Lifecycle demand shares.

After the legislation change, retailers were allowed more flexibility in determining the price. The restraint period was shortened by eight months, and the restraint was now a price band rather than a fixed price. The price band is illustrated by the shaded area in Figure (3). Three changes stand out. Firstly, prices are comparable before and after in the introductory period. The average price lies well within the price band in the shortened restraint period. Despite the retailers new-

found discount discretion under the new restraints, there were hardly any changes in the price patterns in the early phase of the lifecycle. That shows that the new restraints were not materially restrictive of retailers pricing.

The changes in pricing are seen towards the end of the old restraint period in the following year. Here retailers exploited their discount discretion and prices are seen to fall earlier. That shows that RPM was mainly effective in keeping prices high towards the end of the restraint period. It also shows that the pricing incentives changes over the course of the lifecycle, a clear sign of dynamics. Thirdly, prices fall to about the same level at the clearance sale. The industry is not reaching new lower valuation consumer groups by dropping prices deeper than before. The same consumer groups are served, but at different prices over the course of the lifecycle.

The lack of changes in price patterns in the introductory period leaves little explanatory power to the standard rationales for RPM of double marginalisation and complementary service provision. If double marginalisation had been a first order issue, the prices in the introductory period would be expected to increase, yet they largely stay put. As for complementary service provision, since prices in the introductory period hardly drops either, it does not seem to be strong competition between retailers. In sum, the changes in the pricing patterns do not seem well explained by the classic RPM motivations.

Table 2: Changes to introductory prices.

Variable	Coeff.	Std. Err.	Coeff.	Std. Err.
post	4.50	2.25	5.80	1.88
constant	211.79	1.95	228.62	4.60
genre ffx		No		Yes
N		16605		

Figure (4) plots the corresponding shares of sales over the lifecycle. Note that prices in the introductory period were comparable before and after the legislation change. The changes were towards the end of the following year, when the restraints were completely lifted. Demand shifts from comparable early phase prices towards future lower prices. The evidence is consistent with forward looking consumers. Expecting lower future prices following the legislation, consumers are more willing to wait for a future discount at otherwise comparable prices after the deregulation. The change in price strategies and the demand response after the deregulation also goes to show that price skimming falters in the absence of RPM. More demand is served at lower prices over the course of the lifecycle. The changes in both price and demand patterns are qualitatively stable across years and across genres, see Figures , , and in the Appendix for some robustness checks.

The reduced form evidence allows us to ballpark the returns to commitment. The shift in demand is on the order of 10%, the change in prices is on the order of 15%. In total, a ballpark estimate of the profit loss is then 1.5% of total profits before the legislation change. Though not exact, it provides a point of departure for the later quantification exercises.

6 Model setup

The goal of the empirical analysis is quantify the commitment effect of RPM. While RPM provides commitment to future price, it also shuts down competition between retailers. To separate the effect of commitment from effect retailer competition, I turn to counterfactual exercises. The modelling approach follows two steps. In the first step, substitution patterns along the horizontal and inter temporal dimension are estimated from the sales data. In the second step, counterfactual vertical contracts with varying levels of commitment are evaluated in a dynamic oligopoly

model at the estimated parameters. Differences in the profitability between the contracts serve as a measure of the value of commitment.

The oligopoly model has a forward-looking demand side and a forward looking supply side. Retailers set prices taking into account the impact of their prices on both the current and the future demand. The retailers price setting may however be restricted by the vertical contracts. Consumers make purchase decisions considering the current prices and their beliefs about future prices. The beliefs of consumers will depend on the vertical contracts. For instance, consumers beliefs about the future prices depend on whether the industry uses RPM or not. Both consumers and retailers form beliefs over the evolution of prices and demand that are rational in the sense of being consistent with the distribution of the prices and demand in equilibrium.

6.1 Demand specification

The demand is modelled along the lines of Arcidiacono & Ellickson (2011) and is a fairly conventional adoption model. Demand is represented by a finite horizon discrete choice model with a discrete type space. A finite horizon is used as the number of periods is small, it allows for non-stationary policy functions, the typical life cycle is finite, and that it allows for a convenient, yet flexible estimation of the state transition process.

6.2 Consumers, preferences and choices

A key assumption made about substitution is that each title is assumed to be an independent market. The relevant dimensions of substitution are then taken to be between retailers and over time for a given title. This is a strong and restrictive assumption on substitution, but it allows me to parsimoniously focus on substitution between retailers and over time. Allowing for substitution between titles leads to a high dimensional state space that frustrates computation of value functions. Some powerful simplifying assumptions have been developed, like the inclusive value sufficiency assumption in Melnikov (2013), yet these also imply rather strong restrictions on the state transition process. Assuming away substitution between titles allows a more flexible specification of the state transition process within titles, across retailers, which is a central focus of this paper.

There are two types of consumers, a high type h and a low type l , with shares $w^h, 1 - w^h$ of the population, respectively. The lifecycle of a title has T periods. A title is introduced for sale prior to the first period, at which point the share of consumers of each type in the market are the population shares. In the first period, a consumer chooses whether to buy the title or not from one of the J retailers. If he buys, he leaves the market never to return. If he does not buy, the same choice set presents itself in the next period. After T periods, the market closes, and no further choices are made.

A consumer observes the prices $p_t = (p_{1t}, \dots, p_{Jt})$ and exogenous states $x_t = (x_{1t}, \dots, x_{Jt})$ at the start of period t . The exogenous states include deterministic functions of time like seasons, a taste for novelty, and the remaining periods of the lifecycle. They also includes various time invariant characteristics like retailer fixed effects and title specific characteristics, like weight and pages. Some elements of x are hence equal across retailers, and some are time invariant. The exogenous states also allow for retailer and time varying, exogenous demand shocks $\xi_{jt} \in \mathbb{R}$. A full description of the variables that enter the empirical model is given in Section (10). At the start of each period, each consumer privately learns a vector of *iid* EV1 distributed shocks $\epsilon_{it} = (\epsilon_{0t}, \dots, \epsilon_{Jt})$. The utilities are parametrised by vector $\gamma^h = (\gamma_x, \gamma_p)$.

A consumer i of type h who is in the market at time t has current period util-

ities

$$u_{ijt}^h(x_t, p_t, \epsilon_{ijt}) = \begin{cases} \gamma_p^h p_{jt} + x_{jt} \gamma_x^h + \epsilon_{ijt} & \text{if buys from retailer } j \in 1, \dots, J \\ \epsilon_{i0t} & \text{if does not buy } j = 0 \end{cases}$$

where the mean utility from not buying is normalised to zero, and where γ are utility parameters common to the consumer type. The utilities of type l are defined similarly.

Consumers are forward looking expected utility maximisers who discount future utility by factor $\beta \in (0, 1)$. Consumers form expectations F^d over the future states of the market, notably including prices. The consumers choices are characterised by the choice specific value functions $v_{jt}^h : X \times \mathbb{R}_+^J \rightarrow \mathbb{R}$ that measure the expected life time utility of a consumer of type h , conditional on choosing $j \in \{0, \dots, J\}$ in period t . The choice specific value functions in period t are

$$v_{jt}^h(x_t, p_t | \gamma) = x_t \gamma_x^h + \gamma_p^h p_{jt} \text{ if } j = 1, \dots, J$$

if he buys from retailer j , which is just the per-period utility. Alternatively, the current pay-off can be thought of as the present value of lifetime consumption of the book.

A consumer that chooses not to buy, gets a current pay-off of zero, and the expected value of making the optimal choice in the next period when the same choice set presents itself again.

$$v_{j0}(x_t, p_t | \gamma) = 0 + \beta \iint \left(\max_{j=0, \dots, J} v_{jt+1}(x_{t+1}, p_{t+1} | \gamma) + \epsilon_{jt+1} \right) \times dF^d(x_{t+1}, p_{t+1} | x_t, p_t) dF^d(\epsilon)$$

The expectations are taken both over the states of the exogenous states x_t , over the retail prices p_t , and over the independent private shocks ϵ .

In the terminal period, $v_{0T}(x_t, p_t | \gamma) = 0$. Rolling back one period

$$v_{0T-1}(x_{T-1}, p_{T-1} | \gamma) = \beta \iint \left(\max_{j \in \{0, \dots, J\}} v_{jT}(x_T, p_T | \gamma) + \epsilon_{jT} \right) \times dF^d(x_T, p_T | x_{T-1}, p_{T-1}) dF^d(\epsilon)$$

The choice specific value functions are similarly defined up until period $t = 1$ and are solved by backwards recursion conditional on γ .

Consumers choice probabilities, prior to learning ϵ , can be defined in terms of the conditional value functions. Define the indicator

$$d_{ijt}^h = \begin{cases} 1 & \text{if consumer } i \text{ chooses to buy from } j \text{ in period } t \\ 0 & \text{otherwise} \end{cases}$$

Conditional on still being in the market at time t , the consumer's choice probability is

$$Pr(d_{ijt}^h = 1 | x_t, p_t) = \frac{\exp(v_{jt}(x_t, p_t | \gamma))}{\sum_{k=0}^J \exp(v_{kt}(x_t, p_t | \gamma))},$$

which completes the individual choice model.

6.3 Aggregate demand

The aggregate demand is simply the individual choice probabilities summed over all consumers in the market. Suppressing the dependence on γ , the demand shares of type h in period t are

$$D_{jt}^h(R_t^h, x_t, p_t) = R_t^h Pr(d_{ijt}^h = 1 | x_t, p_t)$$

where R_t^h is the residual share of consumers left in the market. Starting from $R_1^h = 1$, the residual demand is defined recursively

$$R_{t+1}^h(x_t, p_t) = R_t^h Pr(d_{i0t}^h = 1 | x_t, p_t)$$

Collect the market level states in $s_t = (R_t, x_t)$, and rewrite the market shares of type h as

$$D_{jt}^h(s_t, p_t) = R_t^h D_{0t-1}^h(s_{t-1}, p_{t-1})$$

Finally, the aggregate demand shares sum the weighted demand shares over both types

$$D_{jt}(s_t, p_t) = w^h D_{jt}^h(s_t, p_t) + w^l D_{jt}^l(s_t, p_t) \quad (1)$$

The consumer expectations F^d have so far been left unspecified. The expectations in the counterfactuals and in the estimation enter differently. In the demand estimation, consumers expectations are estimated using a flexible reduced form from the observed transitions. The expectations in the model are then implemented with quadrature using the estimated moments from a VAR model. The empirical specification of the transitions are described in more detail in Section (9). In the counterfactuals, consumers expectations will be however be determined in equilibrium. More detail is given in Section (6).

7 Supply side

The market for a title has a supply side with one vertical unit that consists of a publisher and a set of J retailers. The demand side is described in the previous section. The publisher produces copies of a given title at a constant marginal cost of production c . A retailer j sets price $p_{jt} \in \mathbb{R}_+$ in period t . The market level states are $s_t = (R_t, x_t)$, where the residual demands are $R_t \in R \subseteq [0, 1]^2$. The exogenous states are $x_t \in X \subset \mathbb{R}^K$. The market level states are commonly observed by the retailers and the consumers at the start of each period. Consumers are assumed to be atomistic, whose decisions individually do not impact the market outcome. The retailers per-period profits are given as

$$\pi_j(s_t, p_t) = (p_{jt} - c) D_{jt}(s_t, p_{jt}, p_{-jt} s)$$

The retailers discount profits by factor ρ , which is considered the cost of capital. According to the World Bank, the mean real interest rate in Norway from 2004 to 2007 was -3%, and the corresponding inflation target was 2.5% in Norway. As an approximation, ρ is set to 0.999.

7.1 Assessing coordination by contracts

The contracts vary with the levels of commitment and the level of coordination of prices across retailers they offer.

- Full commitment *FC*. The vertical unit is endowed with commitment power and coordinates prices both across retailers and over time using RPM.
- Horizontal coordination *HC*. The vertical unit coordinates the prices across retailers using RPM, but has no commitment power.
- No coordination *NC*. The publisher supplies books to the retailers at the marginal cost of production. The retailers engage in dynamic price competition. Serves as a benchmark.

Finally, the time limited RPM contract used in the Norwegian book industry is compared to a set of counterfactual vertical contracts.

The value of coordinating prices horizontally is measured as the difference in the vertical units profits between *HC* and the *NC* contract. The value of commitment beyond horizontal coordination is measured as the difference between the profits of the *FC* and the *HC* contract.

7.2 Full commitment (*FC*)

The vertical unit is endowed with commitment power in the *FC* contract. The timing of moves in the *FC* contract is illustrated in Figure (5). Prior to period 1, the retailers announce a price path $\mathbf{p}' = (p'_{1t}, \dots, p'_{Jt})_{t=1}^T$. The price strategies depend only on time, a subset of the states, and not on all the pay-off relevant states of the market. Consumers know the retailers will price along the announced \mathbf{p}^{FC} price path. The announcement therefore fixes consumers expectations of future prices. The retailers have an incentive to discount the fixed price path over the course of the lifecycle as high valuation consumers clear out of the market. Commitment by assumption prevents the retailers from doing so.

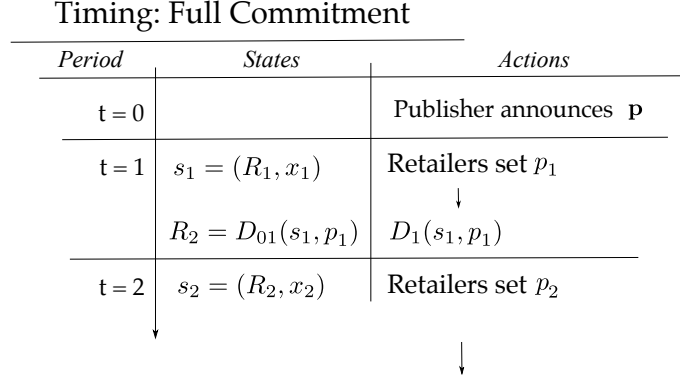


Figure 5

A *Full Commitment* price path \mathbf{p}^{FC}

$$\mathbf{p}^{FC} = \arg \max_{\mathbf{p}} \sum_{t=1}^T \rho^{t-1} \int \pi_t(s_t, p_t) dF^r(s_t | s_{t-1}, p_{t-1})$$

and a *Full Commitment* rational expectations equilibrium is a fixed point such that

$$p_t = p_t^{FC}, D_t = D_t(x_t, p_t), F^d = F^r = F, R_{t+1} = D_{0t}(x_t, p_t)$$

for all $t = 1, \dots, T$

7.3 Contracts without commitment

In the contracts *HC* and *NC* without commitment, prices are set in every period, taking the state of the market into account. The market evolves as a two-stage game within each period, illustrated in Figure (6). First, having observed the current market states s_t , all retailers set prices p_{jt} simultaneously at the start of the period. Consumers observe the prices p_t and demand is realised. All consumers who buy exit the market, and leaves the share of consumer who choose to wait as the next period residual demand R_{t+1} . In period $t + 1$, the same process repeats itself until the last period T , when the market for the title ends.

In the *HC* counterfactual, the publisher decides the prices for all retailers in every period. The price strategies for each retailer is a state contingent price rule and is represented by $\sigma_{jt} : R \times X \rightarrow \mathbb{R}_+$. A strategy profile $\boldsymbol{\sigma}_t = (\sigma_{1t}, \dots, \sigma_{Jt})$ is a collection of retailer specific strategies. Attention is restricted in the following to pure Markov strategies. An MPE is a refinement of sub game perfect equilibrium concept where strategies depend only on the current and pay-off relevant states. An MPE allows for a parsimonious parametrisation and a relatively low dimensional state space.

Timing: Horizontal Coordination

Period	States	Actions
t	$s_t = (R_t, x_t)$	Publisher sets p_t ↓
	$R_{t+1} = D_{0t}(s_t, p_t)$	$D_t(s_t, p_t)$
t + 1	$s_{t+1} = (R_{t+1}, x_{t+1})$	Publisher sets $p_{t+1} \dots$ ↓

Figure 6

The joint profits of the vertical unit are given as

$$\pi(s_t, p_t) = \sum_{j=1}^J \pi_j(s_t, p_t)$$

The vertical unit has no commitment power, so prices are re-optimized in every period. A strategy profile $\sigma_t^{HC} = (\sigma_{1t}^{HC}, \dots, \sigma_{Jt}^{HC})$ solves

$$\Pi_t(s_t, \sigma_t^{HC}) = \max_{p_t \in \mathbb{R}_+^J} \pi(s_t, p_t) + \rho \int \Pi_{t+1}(s_{t+1}, \sigma_{t+1}^{HC}) dF^r(s_{t+1} | s_t, p_t)$$

for all $t = 1, \dots, T$, for all $s_t \in S$.

A rational expectations *HC* Markov Perfect Equilibrium is a fixed point such that

$$p_t = \sigma_t^{HC}(s_t), D_t = D_t(s_t, p_t), F^d = F^r = F, R_{t+1} = D_{0t}(s_t, p_t)$$

for all $t = 1, \dots, T$. The first restriction requires retailers to price according to the price strategies that maximise the joint profits, without commitment, subject to the retailers forecasts of future states and demand. The second restriction requires consumers demand to from expected utility maximisation, subject to the consumer forecasts over the market evolution, notably including prices. The third restriction requires that both consumers and retailers forecasts of future prices and residual demand coincide with the distribution their respective transitions in equilibrium. Importantly, consumers know the retailers price strategies. Without commitment to future prices, consumers find future discounts more likely, which makes a current demand less likely. Finally, the composition of demand evolves according to the specified law of motion.

7.4 No coordination (*NC*)

The *NC* counterfactual is a benchmark contract that represents the pricing of a vertical unit that neither coordinates over time nor across retailers. Publishers supply retailers books at the constant marginal cost of production c . The timing of the market (7). The retailers set prices in every period to maximise their individual expected present value profits, taking into account their rival retailers pricing and the impact of their current prices on future demand. The wholesale price is set equal to c , equal across retailers within the vertical unit. Attention is again restricted to Markov Perfect Equilibriums (MPE).

A *No-Coordination* strategy profile $\sigma_t^{NC} = (\sigma_{1t}^{NC}, \dots, \sigma_{Jt}^{NC})$ is a Markov Perfect Equilibrium if

$$\begin{aligned} \Pi_{jt}(s_t, \sigma_t^{NC}) &\geq \Pi_{jt}(s_t, \sigma_{jt}, \sigma_{-jt}^{NC}) \\ &= \pi_j(s_t, \sigma_{jt}, \sigma_{-jt}^{NC}) \\ &\quad + \rho \int \Pi_{jt+1}(s_{t+1}, \sigma_{jt+1}, \sigma_{-jt+1}^{NC}) dF^r(s_{t+1} | s_t, \sigma_{jt}, \sigma_{-jt}^{NC}) \end{aligned}$$

Timing: No Coordination




Figure 7

for all alternative Markov σ_{jt} , for all $j = 1, \dots, J$, for all $t = 1, \dots, T$, for all $s \in S$.

A rational expectations *No Coordination* equilibrium is now a fixed point that satisfy the equilibrium restrictions.

$$p_t = \sigma_t^{NC}(s_t), D_t = D_t(s_t, p_t), F^d = F^r = F, R_{t+1} = D_{0t}(s_t, p_t),$$

for all $t = 1, \dots, T$. Consumers know that retailers compete in prices, and that affects their expectations of future prices through F^d , which in equilibrium will be correct.

7.5 Time Limited RPM regime

In line with the Trade Agreement, the publisher sets a fixed price \bar{p} , constrained equal across retailers, for a restraint period of τ periods. Following the expiration of the restraint period, retailers set prices freely. The price regime combines the full commitment *FC* strategy and the rivalrous retailer price strategies in *NC*. The publisher set a a fixed price p^{FX} in the restraint period, equal across retailers and periods, taking into account that retailers will employ price strategies σ^{NC} when the restraints are lifted and forwards. The fixed price p^{FX} for the first τ periods is characterised by

$$p^{FX} = \arg \max_{p \in \mathbb{R}_+} \sum_{t=1}^{\tau} \rho^{t-1} \int \pi_t(s_t, p) dF^r(s_t | s_{t-1}, p) + \sum_{j=1}^J \sum_{t=\tau+1}^T \rho^{t-1} \int \pi_j(s_t, \sigma_t^{NC}) dF^r(s_{t+1} | s_t, \sigma_t^{NC})$$

where again commitment is by assumption. The last $T - \tau$ period prices of the strategy are the p^{NC} strategies from above. The wholesale prices are set equal to the marginal cost of production. A typical price path is then

$$\mathbf{p}_j^{TL} = (p^{FX}, \dots, p^{FX}, \sigma_{j\tau+1}^{NC}, \dots, \sigma_{jT}^{NC})'$$

7.6 Existence and multiplicity of equilibriums

I do not have existence theorems, but I find pure strategy equilibriums at the empirically relevant parameters. Multiplicity of equilibria is almost certain given the non-linearity of the model. Yet, preliminary sensitivity tests of the calculated equilibria to perturbations of the parameters local to convergence do not reveal issues of discontinuous jumps in policy functions. I have found reaction curves to intersect locally at convergence. Further discussion on solution methods are relayed to Section (11) on results.

8 Estimation

This section discusses the identification the substitution patterns informally before turning to the estimation routine.

8.1 Identification

The discount factor which in part determines the inter temporal substitution is usually not identified in dynamic models, i.e. there is not a unique vector of discount factor β and utility parameters γ that can rationalise the data. Estimation usually proceeds conditional on a discount factor that is fixed prior to estimation. As the value of commitment crucially depends on the discount factor, fixing the discount factor in some sense assumes the result. That is a source of concern. The price variation generated by the legislation change allows the discount factor to be identified.

In the ideal experiment, the discount factor is identified by measuring how the current demand responds to changes in future utility, holding the current utility fixed (Magnac & Thesmar (2002)). For a myopic consumer, only the current price matters, whereas a forward looking consumer responds to changes to future prices. The identifying experiment is arguably in the data. From Figure (3), prices in the year of publication are comparable both before and after the deregulation. Yet demand in the introductory period shifts towards *future* lower prices after the deregulation. The price variation generated by the deregulation is close to the identifying experiment, and the variation that I use to identify the discount factor.

Having been in place since the early 1960s, the fixed price model was widely known among the book buyers. Consumers could be reasonably sure there would be no discounts on the introductory price until the expiration of the restraint period. These expectations are likely to have changed after the deregulation as price paths are seen to materially change. The expectations F^b , F^a before and after the legislation change is estimated from the data in Section (9) .

The standard endogeneity problem in demand estimation arises if retailers rather than randomly vary the prices, set them in response to changes in demand. There is then something observable to consumers and to retailers that affects both the demand and the price incentives that is unobserved in the data. The usual solution is to find instrument variables that are correlated with retailers price incentives, but do not directly affect demand itself. Standard sources of instruments include marginal cost shifters (Working (1927)), variation in the density of the product space (BLP (1995)) and geographical price variation (Hausman (1996), Nevo (2001)).¹⁴ None of these standard sources are readily available in the current application. For books, marginal costs are likely close to constant over the lifecycle of a title, non-price product characteristics are mostly time-invariant and there is no geographical variation in the data. Since most of the regular instruments are likely weak, questionable or unavailable, I let prices instrument for themselves.

To ameliorate the endogeneity issue in the absence of instruments, the approach is to find data that can proxy for the unobservable demand shocks that cause the endogeneity problem. Though at least partially controlling for demand shocks does not solve the problem, the idea is that it can reduce the problem. Importantly, this approach is not a perfect substitute for valid instruments, but it can hopefully reduce the bias induces by endogeneity. For that purpose, editorial book reviews have been collected and converted to a quality variable. A total of 1823 reviews for the period 2002 to 2007 were collected from the three largest national circulation newspapers (Aftenposten, Dagbladet and VG), in print and on the web.¹⁵

¹⁴In its simplest form, the geographical instruments use that whereas demand shocks may be local, cost shocks may be national. The idea is then to decompose the price variation in a national component and orthogonal local variation. The shared national variation is used as to instrument for prices.

¹⁵The reviews were accessed through Retriever, a comprehensive and proprietary Scandinavian media archive.

About a third of the text reviews were editorially scored by the newspaper with a grade ranging from 1 to 6 in addition to the text review. The rest were editorially unscored. Just the fact that a title is reviewed is positively predictive of sales, but higher scores also predict higher sales.¹⁶ To extract this additional information from the unscored reviews, a simple supervised machine learning algorithm was applied.

The scored reviews were used as a training set to define an algorithm that could predict the unscored reviews. The review samples are very small for text analysis, so a premium was put on simple algorithms and crude scoring that could robustly capture the important feature. Unscored reviews were predicted as either positive or negative. A k Nearest Neighbours-approach with simple majority voting was used, with k set somewhat arbitrarily to 5.¹⁷ Based on the frequency of indicative words in a (stemmed) unscored review, its five closest reviews from the training set in terms of number of shared indicative words were found. Then the vote of the five nearest neighbours was cast. If three or more neighbours were positive, the review was classified as positive, negative otherwise. The scored reviews were converted to the same scale. The algorithm was validated on a set of titles with both editorially scored and unscored reviews.

I use the approach in Nair (2007) to infer the market size on title level from cumulative sales by appealing to the Bass (1969) model of diffusion. The model specifies the market size as a flexible function of the cumulative sales of each title. The market shares are then constructed from observed sales and the estimated market size, see the Appendix for details.

8.2 Specification

To avoid issues with observing market shares of exactly zero, only titles with strictly positive sales at all retailers for consecutive periods since introduction were used for the estimation of utility parameters.

The exogenous state variables x_{jkt} for retailer j for title k in period t are partitioned in deterministic states z_{jkt} and demand shocks ξ_{jkt} . The deterministic states are $z_{jkt} = \{ttl_k, r_j, n_t, ssn_t, rvw_{kt}\}$ where ttl_k is a title fixed effect and r_j is a retailer fixed effect. Following Einav (2007) and Ho et. al. (2012), the utility has a component $\gamma_n n_t$ that represents a taste for novelty, where n_t is a function of time. Seasonal fixed effects are given by ssn_t , and rvw_{kt} are review state variables. Title specific fixed effects are given by γ_k . Discrete type heterogeneity is allowed on valuations and price coefficients, so that high type has valuation parameters γ^h and price parameter γ_p^h . The current period utility is then

$$\begin{aligned}
 u_{ijkt}^h &= \gamma_k + \gamma_j r_j + ssn_t \gamma_{ssn} + \gamma_{rvw} rvw_{kt} + \gamma_n n_t + \gamma_p^h p_{jkt} + \gamma^h + \epsilon_{ijkt} \\
 &= \underbrace{x_{jkt} \gamma}_{\text{mean utility}} + \underbrace{\gamma^h + \gamma_p^h p_{jkt}}_{\text{type specific utility}} + \underbrace{\epsilon_{ijkt}}_{\text{individual specific utility}}
 \end{aligned}$$

where the last line partitions the current period utility mean, a type specific $\gamma_p^h p_{jkt} + \gamma^h$ and an individual specific utility component. The specification for type l is defined similarly.

8.3 Estimation routine

The estimation proceeds sequentially. First the consumers expectations F^d are estimated in reduced form from the data. The estimated transition process then enters the demand specification. The estimation routine is described in reverse order. I first describe the estimation of the utility parameters γ , which is fairly standard, and then outline the estimation approach to the discount factor β , which

¹⁶See the Appendix for results

¹⁷See Hastie, Tibshirani & Friedman (2008) for an exposition.

is somewhat less standard, before I turn to the estimation of the expectations.

The structural parameters β, γ are estimated using GMM.¹⁸ The estimation is implemented in a two-step routine where the optimal weighting matrix W is estimated in the first step. The BLP (1995) routine is used for estimation. Though there are no exclusion restrictions, the BLP mean utility inversion is used as a computationally convenient, if inefficient, way to concentrate out the title fixed effects and the shared mean parameters. Integrating out ϵ , the only errors left in the demand function to rationalize the data are the unobserved demand shocks ξ . The mean utility across groups is collected as

$$\tilde{\gamma}_{jkt} = \gamma_k + \gamma_j r_j + ssn_t \gamma_{ssn} + \gamma_{rvw} r v w_{kt} + \gamma_n n_t + \xi_{jkt} = \tilde{x}_{jkt} \gamma_{\tilde{x}} + \xi_{jkt}$$

The routine proceeds using the inner/outer loop routine of BLP (1995). Conditional on the γ parameters, the shared component $\tilde{\delta}_{jkt}$ is inverted out from the restriction

$$\tilde{\gamma}_{jkt} : D_{jkt} = D_{jkt}(\tilde{\gamma}_{jkt}, \gamma)$$

where D_{jkt} is the observed demand and $D_{jkt}(\tilde{\gamma}_{jkt}, \gamma)$ is the predicted demand. The parameters of the shared utility component δ are concentrated out by a linear regression of $\tilde{\gamma}_{jkt} = \tilde{x}_{jkt} \gamma_{\tilde{x}} + \xi_{jkt}$, where ξ_{jkt} is assumed *iid*. The inversion itself is implemented by the BLP contraction mapping

$$\tilde{\gamma}^{it+1} = \tilde{\gamma}^{it} + \ln(D_{jkt}) - \ln(D_{jkt}(\tilde{\gamma}^{it}, \gamma))$$

The contraction is run title-by-title until convergence.¹⁹

In the outer loop, the residuals $\hat{\xi}_{jkt}$ are interacted with the remaining covariates x, p under the mean independence assumption $E[\xi|x, p] = 0$. The remaining γ^h, γ^l parameters are found in an outer loop search as

$$\hat{\gamma}^f = \arg \min_{\gamma} m^f(\gamma) W m^f(\gamma)' \quad (2)$$

where $m^f(\gamma) = \xi(\gamma)' X$, and W is the optimal weighting matrix.

The empirical strategy for the estimation of the discount factor is less standard. The idea is to add a moment that lets the discount factor explain the changes in the introductory period demand before and after the legislation change at otherwise comparable prices. Denote the observed aggregate demand before and after the change D_t^b, D_t^a , respectively. The identifying variation comes from the shift in consumers expectations F^b, F^a induced by the change of price strategies following the legislation change. Let L^b, L^a be the number of titles sold before and after, respectively. The left-hand side is the difference in demand before and after in the same periods of the lifecycle. The right hand side is the theoretical demand that lets the changes in expectations explain the difference in demand at otherwise comparable prices. The moments are then

$$\frac{1}{L^b} \sum_t D_{lt}^b - \frac{1}{L^a} D_{lt}^a = \frac{1}{L^b} \sum_L D_{lt}(p^b, F_b) - \frac{1}{L^a} D_{lt}(p^a, F_a)$$

for each $t = 1, \dots, T$ over the lifecycle. These moments are added to (2). The expectations F^b, F^a are estimated from the data as described below.

9 State transitions and expectations

The empirical specification of the state transition process is a simple linear-in-parameters Markov process with additive shocks. The specification reflects that

¹⁸Software matters: Both the estimation and the counterfactuals use SNOPT and automatic differentiation extensively. A special thanks to Mike Saunders for making the software available for free and similar thanks to Anders Goran at TomLab for at request adding differentiation modules to MAD.

¹⁹Following Dube, Fox & Su (2012), with tight tolerances.

consumers may form beliefs about the future prices taking into account the current prices at all retailers and factors such as what phase of the lifecycle the title is in and what genre the title is. Time is modelled with a flexible specification, so the process is Markov conditional on life cycle fixed effects. The specification allows consumers to form expectations from observing variation in price paths across titles, and allow expectations to vary with retailers and different genres of books.

Let the consumer states be partitioned in the stochastic and deterministic processes $s_t = (p_t, z_t)'$, respectively.

$$s_{t+1} = \Theta s_t + \eta_t$$

where $\sigma_{\eta_z} = 0$, by definition of deterministic z . Each retailer is allowed individual coefficients Θ_j . The parameters of the specification can be partitioned

$$\Theta = \begin{bmatrix} \theta_{pp} & \theta_{pz} \\ 0 & \theta_{zz} \end{bmatrix}$$

where the lower left zero matrix follows from the exogeneity of z . By forward iteration, we can then write

$$s_{t+r} = \Theta^r s_t + \sum_{\tau=1}^r \Theta^{r-\tau} \eta_\tau$$

In line with the Markov Perfect Equilibrium concept used in the counterfactuals, the current state s_t is assumed to summarise all the payoff relevant information to the consumer and hence doubles as the information set.

The shocks η are assumed mean zero multivariate normal and serially uncorrelated

$$E[\eta_t | s_t] = 0 \text{ for all } t \quad (3)$$

$$E[\eta_t \eta_t' | s_t] = \Sigma \text{ for all } t \quad (4)$$

$$E[\eta_t \eta_{t+r} | s_t] = 0 \text{ for all } t \neq r \quad (5)$$

The assumptions on η jointly define a martingale difference sequence adapted to the information set s_t . The expectations conditional on s_t are now simple functions of s_t itself.

$$\begin{aligned} E[s_{t+r} | s_t] &= \Theta^r s_t + \sum_{\tau=1}^r \Theta^{r-\tau} E[\eta_{t+\tau} | s_t] \\ &= \Theta^r s_t \end{aligned}$$

since $E[\eta_{t+\tau} | s_t] = 0$ for $\tau \geq 1$. Under assumptions (3)-(5), the second moment is

$$\begin{aligned} V[s_{t+r} | s_t] &= V \left[\sum_{\tau=1}^r \Theta^{r-\tau} \eta_{t+\tau} | s_t \right] \\ &= \sum_{\tau=1}^r \Theta^{r-\tau} \Sigma \Theta'^{r-\tau} \end{aligned}$$

The specification is estimated by a FGLS for each group. The process is stable at $\hat{\Theta}$ with the characteristic roots of the price parameters all being of modulus less than one. The model fits quite well with $R^2 = 0.96$. See the Appendix for further details.

Following the arguments in Skrainka & Judd (2011), quadrature is used to numerically integrate out the expectations in the demand functions. Normality of η motivates the use of multivariate Hermitian quadrature, see Judd (1998) for an exposition. Normality along with rational expectations also implies that the two first moments of the transition process completely describe the expectations F . The option value integrals

$$v_{0t-1}(s_{t-1}) = \int \left(\max_{j \in \{0, \dots, J\}} v_{jt}(s_t) \right) dF(s_t | s_{t-1})$$

are solved with tensor product bases with five nodes in each dimension.

10 Estimation results

The parameter estimates are given in Table (3). Some patterns emerge. The consumers discount factor β comes out at 0.68, substantially less than the real interest rate which is used as the retailers discount factor. The estimate is small, but in line with Dube et al (2012). There is an estimated overweight of high valuation consumers, the share comes out at three quarters. Seasonality is estimated, with fall being the premium season. There is some variation in retailer fixed effects, reflecting product differentiation. As for the type specific parameters, the high valuation consumers are also less price sensitive.²⁰ By and large, the estimates are mostly significantly different from zero, with some exceptions. The price coefficient of the high type is not, but we will see that the specification has reasonable implications in the counterfactuals.²¹

Table 3: Parameter estimates

Mean utility parameters				
	Coeff	Std Err		
β	0.68	0.25		
w^h	0.74	0.39		
novelty	0.24	0.09		
review	0.88	0.39		
spring	0.26	0.22		
summer	0.39	0.24		
fall	0.88	0.38		
retailer A	0.00	n.a.		
retailer B	0.59	0.48		
retailer C	0.67	0.23		

Type specific parameters				
Parameter	Low		High	
	Coeff	Std Err	Coeff	Std Err
valuation	-0.12	0.23	2.20	0.97
p	-1.38	0.39	-0.63	0.38
N	12796			

Retailer A serves as base level.

11 Counterfactual results

The counterfactuals are calculated at the preliminary demand estimates and at the mean of the exogenous, deterministic states z in the sample. For each of the counterfactual contracts, the price strategies described in Section (6) are calculated along with the corresponding demand and profits. The marginal cost is constant over time and set to 0.5, in line with industry estimates. The equilibrium contracts are calculated for a representative title over the course of a lifecycle of seven periods, starting in spring of the year of introduction. Note that ξ is set equal to zero for now, effectively turning the models into perfect foresight. For ease of exposition, the prices and demand have been averaged over the retailers. The results are plotted in Figure (8).

The FC price path is seen to lie everywhere higher than the HC price path. With coordination and commitment, the vertical unit keeps prices higher over the course of the lifecycle to induce consumers to purchase early. Without commitment, prices unravel and are seen to drop deeper. Note that the model does not generate prices in the FC equilibrium that falls to the same level as in the HC equilibrium. That implies rationing in FC equilibrium, commitment implies that demand is rationed

²⁰The fixed effects specification improves much on an earlier random effects specification that also was specified with substantially more heterogeneity, that had non sensible economic implications.

²¹The estimation is so far conducted with a randomly selected subset of the titles in the full sample to reduce computational costs.

under FC relative to HC. The rationing in FC stands in contrast to the price patterns in the data, where prices fell conspicuously to the same level before and after the legislation change.

In the NC equilibrium, each retailer set prices to maximise individual profits over the lifecycle, without commitment. Two features of the equilibrium stand out. Firstly, the price paths flatten out, the NC equilibrium price paths have substantially lesser slope. Any individual retailer lacks the market power to implement significant price skimming strategies. A retailer that attempts a price skimming strategy like in say HC would see its demand lost to rival retailers responding with lower prices. The loss of demand to rival retailers dominate the returns to charging high valuation consumers higher prices. Secondly, the all-over price level in NC lies about the same level as HC eventually reaches.

The demand patterns for both high and low valuation demand is qualitatively similar in all equilibriums. In FC equilibrium, the distribution of demand shifts to earlier in the lifecycle, where the demand is served at higher prices. The shift in demand to earlier, higher prices is however counteracted by the effect of rationing. The total demand served in the FC equilibrium is smaller. The profit paths are therefore similar. The NC equilibrium demand is materially different. Since both price paths flatten out and the price level drops, consumers make their purchases earlier in the lifecycle. Even though prices are everywhere lower than FC and HC, the NC demand eventually drops lower after about a year as consumers have cleared out of the market.

In terms of profits, the profits of FC and HC are strikingly similar over the course of the lifecycle. The profit increase of the FC contract over the HC contract lies mainly in the first periods. Here FC can sustain higher demand at higher prices since consumers are also expecting future higher prices. Later in the FC lifecycle, the higher FC prices are however counteracted by lower demand. That points to the time-inconsistency issue in FC: As soon as high valuation consumers are served in the FC equilibrium, the vertical unit has an incentive to lower the price to profit on the low valuation demand still in the market. But if it lowered the prices later in the lifecycle, and this was expected by the consumers, it would not achieve the surplus profit in the first period. The total profits from the three contracts are given in (4). Horizontal Coordination improves industry profitability by 15.3% over the No Coordination baseline, Full Commitment by another 2.5%. This is the current estimate of the value of commitment.

Table 4: Counterfactual profits

Contract	Total profits	Percent increase over <i>NC</i>
FC	1.047	17.8
HC	1.024	15.3
NC	0.888	n.a
TL	1.023	15.2

Numerically, the relative price discounts implied in the model are on par with the data. The HC contract generates about 50% discounts at the end of the lifecycle, compared with 45% in the data at the time of the clearance sale. Note that the clearance sale is not explicitly modelled in the counterfactuals. Note also that there is nothing in the model that explains the bounce-back in the data. In levels, the counterfactuals overshoot the price levels in the market, and quite a bit, by about 40%. That points to the valuations being overestimated. Note here that no supply side restrictions were imposed in the estimation to ensure consistency between the pricing in the data and pricing at the estimated parameters. The discrepancy however suggests that the valuations are estimated to high. A second caveat is the rationing implied by the FC pricing contract. That rationing is not found in the data, where on the contrary prices fell to the same level before and

after the legislation change.

Remove RPM, and there is still substantial price discrimination and the data shows that there was still substantial price skimming. That implies that removing RPM did not release strong horizontal competition between retailers. That would have flattened out the price paths. The pattern after the legislation change is consistent with the vertical unit losing commitment, since prices are falling earlier. It is not known what vertical contracts the industry turned to after the legislation. The evidence however suggest that the vertical unit still coordinates the supply chain, for instance using sequentially decreasing wholesale prices to implement declining retail price paths. These are similar to the decreasing wholesale prices Nair (2007) reports anecdotal evidence for in the US video game industry.

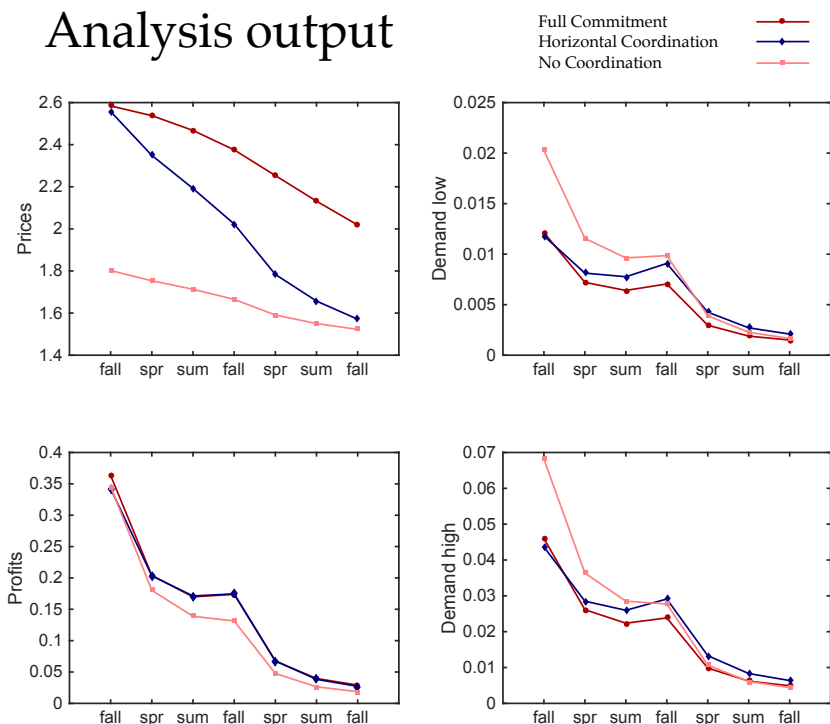


Figure 8

11.1 Performance of time limited RPM

The time limited RPM agreement (TL) after the legislation did not imply fully specifying a price path like in FC. The TL strategy fixed the retail price for a number of periods, followed by retailers pricing the title at discretion after the restraint period expired. The strategies in Section (7.5). The results are plotted in Figure (9) and the profits are reported in Table (4). The profits of the TL strategy are seen to be almost exactly equal to the HC profits. The price and demand patterns however differ substantially. The TL strategy achieves a higher degree of separation of consumers. It has more consumer buying early. On the other hand, it serves more consumers at lower prices later in the lifecycle. Finally, the TL strategy is seen to ration less than the FC strategy. Like in the data, the TL strategy eventually falls to a terminal price level comparable to the HC and NC strategies. The TL profit estimates provide a lower bound in the sense that the wholesale price has been set equal to the marginal cost of production. Optimising the wholesale price, the vertical unit can further improve on profits by coordinating the horizontal coordination in the final periods.

The TL strategy is somewhat heuristic. It fixes the price for a seemingly arbi-

Analysis output

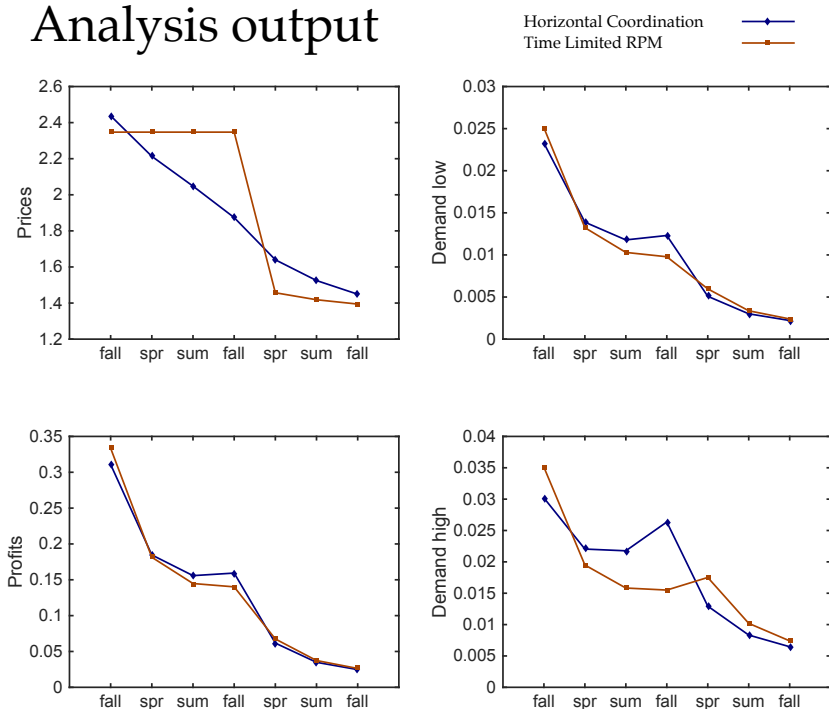


Figure 9

trary number of periods, and it does not discriminate the price between retailers. It finally leaves to retailers to sort out set prices in the final periods, rather than coordinate the horizontal competition. Yet, the heuristic may be quite effective. The TL is simple to understand from a consumer point of view, which is crucial for a commitment strategy. With uncertain demand, the TL strategy also trades off flexibility to adapt pricing to varying demand conditions with commitment to an announced price path. The profits of the TL strategy come close to the profits realised under the HC strategy. It may be that leaving to retailers to adapt the pricing to local demand conditions may be better than letting publishers coordinate the pricing policies centrally at the last stages of the lifecycle. Finally, the TL policy may be a cost effective way of implementing vertical control of titles. Micro managing thousands of titles from a publishers level is a demanding supply chain task with non-trivial transaction costs.

11.2 Caveats

An informal test of fit is to see how well the counterfactual model predicts introductory prices under the Fixed Price counterfactual that mimics the actual price policy. The counterfactual fixed price model currently predicts too high introductory prices. The mean introductory price in the data is about 1.80, whereas the predicted in the model is about 50% higher. The FC equilibrium implies rationing not seen in data, the final price level is about 20% too high.

There are other reasons why the model produce relatively small effects of commitment. The value of commitment in the discrete time models depend crucially on the periodisation.²² In the model, the periodisation follows the data sampling intervals of four months. The periodisation implies a piecewise commitment. Each retailer by assumption is in the model committed to its price for a period of four months. The periodisation is an artefact of the data sampling and does not correspond to any data on how retailers actually price. In reality, the retailers can

²²Stokey (1981) shows that as the length of the period goes to zero, the price of the monopolist goes to cost.

change prices more frequently. In summary, the model therefore likely overstates the profits of a no-commitment strategy. There is a limit to how much can be done to ameliorate the problem without data on the actual frequency of price changes. One simple robustness check is to carve up each four month period in shorter sub periods and allow the retailers to change prices in every sub-period. This robustness check has not been implemented, but the impact is likely to be substantial.

In the current version, I am abstracting away from the idiosyncratic demand shocks ξ in the counterfactuals. Excluding demand shocks inflates the profitability of the commitment strategy over a time-consistent strategy since one advantage of the time-consistent policy is that it can adapt to demand shocks, whereas a fixed price policy can not. For instance, a retailer that learns of unexpectedly high demand of a title may wish to increase its price, or similarly, the retailers may want to clear out titles that turn out to be in low demand by lowering the prices. By excluding demand shocks, the potentially profitable flexibility of the time consistent policy is effectively assumed away. The trade off between commitment and flexibility is important, and it is possible to address within the current framework. The relatively modest estimated returns to commitment suggest that price flexibility may override the profitability of commitment.

12 Discussion

This paper has shown that RPM can improve on dynamic pricing by providing commitment. The Norwegian book industry was regulated by a Trade Agreement which specified an enforcement mechanism to keep its member to the time-limited RPM pricing model. Similar RPM models are used in many European countries. The enforcement mechanism varies, in some countries such as Netherlands and Hungary, the pricing model is enforced by a Trade Agreement. In other countries such as Spain, Portugal, Greece and Mexico, the contracts are enforced by the government.²³ The argument in this paper is that the enforcement mechanisms can provide commitment to future prices and counteract incentives to deviate from the announced price path. The result is that these agreement may improve on price skimming strategies.

Commitment may also stem from less formal sources. For instance, in the US e-book market, the publishers agency model coupled with a Most Favored Nation clause for Apple may have act as a commitment to the RPM pricing model. Furthermore, Apple has built a reputation for rarely discounting their products that may benefit the publishers in the retail market. If consumers believe that e-books on Apples platform would not be discounted, it may have helped the industry to price discriminate. That RPM is so popular in book markets in particular is likely related to the relative ease with which the industry is allowed exemptions. In jurisdictions where RPM is legal, it has become a popular price model in digital markets for other Intellectual Property products like music (iTunes store) and software (AppStore).

Finally, this paper does not argue that commitment is the only effect of RPM, but a complementary effect. The Norwegian book market is not the only European book market that has had RPM abolished. The effects of abolishing RPM has been mixed, see Canoy et al (2005). The mixed evidence suggests that RPM has multiple effects on pricing and demand. In the case of the Norwegian market, the results however show that inter temporal effects are the dominant ones. There is little evidence that avoiding double marginalisation or ensuring retail service provision were important effects of RPM in this market.

²³The International Publishers Association provides an overview at <http://www.internationalpublishers.org/>

13 Summary

This paper studies the role of RPM as part of a price skimming strategy. It starts with the observation that RPM fixes the price path in dynamic markets. By fixing the price path, RPM can act as a commitment device. Comprehensive retail data from a natural experiment in a book market where RPM was deregulated shows a clear impact of RPM on price strategies and demand. Following the deregulation, prices fall earlier over the lifecycle, and demand shifts from early in the lifecycle at high prices to later at lower prices. Fixing the price paths also eliminates horizontal competition between retailers. To separate out the impact of RPM on inter temporal price discrimination, I turn to evaluation of counterfactual vertical contracts. Estimates from a dynamic demand model enter as input to a dynamic supply side model that allows for various levels of price coordination, both across retailers and over time. The estimated value of commitment beyond supply chain coordination is estimated to be on the order of 2.5% of profits, whereas horizontal coordination increases profits by 15% over a benchmark market outcome where pricing is neither coordinated over time nor between retailers.

14 References

- Arcidiacono & Ellickson (2011): "Practical Methods for Estimation of Dynamic Discrete Choice Models", *Annual Review of Economics*, Vol. 3
- Amemiya (1985): "Advanced Econometrics", Harvard University Press
- Berry, Levinsohn & Pakes (1995): "Automobile Prices in Market Equilibrium", *Econometrica*, Vol. 63, No. 4
- Besanko, Dube & Gupta (2005): "Own-Brand and Cross-Brand Retail Pass Through", *Marketing Science*, Vol. 24
- Bulow (1982): "Durable Goods Monopolists," *Journal of Political Economy*, Vol. 90
- Cachon (2003): "Supply Chain Coordination with Contracts", in *Handbooks in Operations Research and Management Science: Supply Chain Management*, eds. Graves and de Kok, North-Holland.
- Chevalier & Goolsbee (2009): "Are Durable Goods Consumers Forward-Looking? Evidence from College Textbooks", *Quarterly Journal of Economics*, Vol. 124
- Coase (1972): "Durability and Monopoly" , *Journal of Law and Economics*, Vol.15
- Conlon (2012): "A Dynamic Model of Prices and Margins in the LCD TV Industry", wp.
- Daljord (2014b): "Dynamic Pass-Through and Integrating Vertical Contracts", wp
- De los Santos & Wildenbeest (2014): "E-book Pricing and Vertical Restraints", wp
- Desai, Koenigsberg & Purohit (2004): "Strategic Decentralization and Channel Coordination", *Quantitative Marketing and Economics*, Vol. 2
- Dube, Fox & Su (2012): "Improving the Numerical Performance of the Static and Dynamic Aggregate Discrete Choice Random Coefficients Demand Estimation", *Econometrica*, Vol. 80
- Dube, Hitsch & Jindal (2012): "The Joint Identification of Utility and Discount Functions From Stated Choice Data: An Application to Durable Goods Adoption",

NBER NBER Working Paper No. 18393

Einav (2007): "Seasonality in the U.S. Motion Picture Industry", *The RAND Journal of Economics*, Vol. 38

Farell & Shapiro (2010): "Antitrust Evaluation of Horizontal Mergers: An Economic Alternative to Market Definition", *B.E. Journal of Theoretical Economics, Policies and Perspectives*.

Goettler & Gordon (2011): "Does AMD Spur Intel to Innovate More?", *Journal of Political Economy*, Vol. 119

Gowrisankaran & Rysman (2009): "Dynamics of Consumer Demand for New Durable Goods", *Journal of Political Economy*, Vol. 120

Gul (1985): "Non-Cooperative Collusion in Durable Goods Oligopoly", *The RAND Journal of Economics*, Vol. 18

Gundlach (2010): "Overview and contents of the special issue: Antitrust analysis of resale price maintenance after Leegin", *The Antitrust Bulletin*, Vol. 55, No 1

Hastie, Tibshirani & Friedman (2008): *The Elements of Statistical Learning: Data mining, inference, and prediction*, *Springer series in statistics*

Hausman (1996): "Valuation of New Goods under Perfect and Imperfect Competition", in *The Economics of New Goods*, eds. Timothy F. Bresnahan and Robert J. Gordon

Ho, Ho & J. Mortimer (2012), "The Use of Full-line Forcing Contracts in the Video Rental Industry," *American Economic Review*, Vol. 102

Ioannou, Mortimer & Mortimer (2011), "The Effects of Capacity on Sales Under Alternative Vertical Contracts," *Journal of Industrial Economics*, Vol. 59

Ippolito (1991): "Resale Price Maintenance: Evidence from Litigation", *Journal of Law and Economics*, Vol. 34

Judd (1998): "Numerical Methods in Economics", The MIT Press. Magnac & Thesmar (2002): "Identifying Dynamic Discrete Decision Processes", *Econometrica*, Vol. 70

Melnikov (2013): "Demand for Differentiated Durable Products: The Case of the US Computer Printer Market", *Economic Inquiry*, Vol. 51

Newey & McFadden (1999): "Large Sample Estimation and Hypothesis Testing", In *Handbook of Econometrics*. Vol. 4. Eds. McFadden & Engle.

McKay & Smith (2013): "The Empirical Effects of Minimum Resale Price Maintenance on Prices and Output", working paper

Mortimer (2008), "Vertical Contracts in the Video Rental Industry," *The Review of Economic Studies*, Vol. 75, pp. 165-199.

Nevo (2001): "Measuring Market Power in the Ready-to-Eat Cereal Industry", *Econometrica*, Vol. 69

Nair (2007): "Intertemporal Price Discrimination with Forward-looking Consumers: Application to the US market for console video games", *Quantitative Marketing*

and *Economics*, Vol. 5

Oery (2013): "Consumers on a Leash: Advertising Sales and Intertemporal Price Discrimination", wp

Overstreet (1983): "Resale Price Maintenance: Economic Theories and Empirical Evidence", Bureau of Economics Staff Report, Federal Trade Commission

Rao (2014): "Online Content Pricing: Purchase and Rental Pricing", wp.

Rao & Srinivasan (1995): "Why are Royalties Higher in Service Type Franchises?", *Journal of Economics and Management Strategy*, Vol. 4.

Rust (1987): "Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher", *Econometrica*, Vol. 55

Shapiro (1996): "Mergers with Differentiated Products, *Antitrust*, Spring

Spengler (1950): Vertical Integration and Antitrust Policy, *Journal of Political Economy*, Vol. 50

Skrainka & Judd (2011): "High Performance Quadrature Rules: How Numerical Integration Affects a Popular Model of Product Differentiation", wp

Stiglitz & Rey (1988): "Vertical Restraints and Producers Competition", *European Economic Review*, Vol. 32

Stokey (1981): "Rational Expectations and Durable Goods Pricing", *The Bell Journal of Economics*, 1981

Su & Judd (2012): "Constrained Optimization Approaches to Estimation of Structural Models", *Econometrica*, Vol. 80

Telser (1960): "Why Should Manufacturers Want Fair Trade 1" , *Journal of Law and Economics*, Vol. 3

Villas-Boas (1998): "Product Line Design for a Distribution Channel", *Marketing Science*, Vol. 17

Working (1927): "What do Statistical "Demand Curves" show?", *Quarterly Journal of Economics*, Vol. 41

A Calculating market shares

To calculate the market size of title k , run the regression

$$q_{kt} = a_k + bQ_{kt} + cQ_{kt}^2 + \epsilon_{kt}$$

with q aggregated across retailers for each period. The market size is then $\hat{M}_k = \frac{\hat{a}}{x_k}$, where x_k is the positive root of the quadratic equation $x_k^2 + bx_k + ac = 0$.

A.1 Changes in release dates

The release dates are from the Bokbasen database. Bokbasen is a company that supplies the industry with information on titles and is used for logistical purposes across retailers and publishers. The release dates are recorded as the time the publishers register a title in the logistical database. There are known to be some error in the release date data, but data as is show no discernible change in release dates before and after the deregulation. Taking the list of titles and regressing the tertile of release on a dummy for the deregulation, the results are given in

Table (??). There are hardly changes in the release dates. The changes in release dates are measured in shares of a tertial and are close to zero. The median title is released in the second tertial, both before and after the deregulation.

Table 5: Change in introductory price levels before and after the deregulation

Variable	Coeff.	Std. Err.	Coeff	Std. Err.
post	0.003	0.014	-0.003	0.015
Intercept	2.063	0.013	2.082	0.036
Fixed effects	No		Yes	
N	16605			

B Estimation of the state transitions

The state transition parameters of the price process $\Theta_p = (\theta_{pp}, \theta_{pz})$ and Σ are assumed known by the consumers, but unknown to the econometrician and is estimated from the data. The parameters of the deterministic process θ_{zz} are known. The transition process is equivalent to SUR specification with lagged prices and is estimated by a standard Feasible Generalized Least Squares approach.

Stacking the price equations for each title l , we get

$$\begin{pmatrix} \mathbf{p}_1 \\ \vdots \\ \mathbf{p}_L \end{pmatrix} = \begin{pmatrix} \theta_{pp}\mathbf{1p}_1 + \theta_{pz}\mathbf{z}_1 \\ \vdots \\ \theta_{pp}\mathbf{1p}_L + \theta_{pz}\mathbf{z}_L \end{pmatrix} + \begin{pmatrix} \eta_1 \\ \vdots \\ \eta_L \end{pmatrix} \quad (6)$$

where $\mathbf{1p}$ are lagged prices. We now have standard linear system of equations that by reorganising the parameters Θ can be written as

$$\mathbf{p} = \mathbf{S}\Theta_{FGLS} + \eta \quad (7)$$

The model is estimated with a two-step procedure. In the first stage, OLS is run on (7). The covariance $\hat{\Omega}$ is estimated as

$$\hat{\Omega} = I_{LT} \otimes \hat{\Sigma}$$

where $\hat{\Sigma} = \frac{1}{LT}\hat{\eta}\hat{\eta}'$ is a consistent estimator of covariance from the first stage residuals of OLS. The FGLS estimator is now the familiar

$$\Theta_{FGLS} = \left(\mathbf{S}'\hat{\Omega}^{-1}\mathbf{S}\right)^{-1}\mathbf{S}'\hat{\Omega}^{-1}\mathbf{p}$$

which completes the estimation of the price transition process. Note that the transition process is estimated independently of the demand system. By standard arguments, the estimated errors $\hat{\eta}$ are asymptotically normal with $\hat{\Sigma}$ a consistent estimator of the covariance.

The estimated moments are given in Table (6) and (7).

The eigenvalues of the matrix of the price coefficients having eigenvalues 0.02, 0.04, 0.09, each well below 1, for the Fiction titles, the Non-Fiction eigenvalues are similar. The estimated coefficients show substantial dynamics and correlations across retailers and over time within titles. The second moments Σ are given in Table (7) in correlation form. After controlling for period fixed effects and lagged prices, there are still substantive positive correlations in pricing across retailers. A multivariate portmanteau test reveal mild signs of serial correlation in η .

C Review text analysis

A total of 1823 reviews were crawled for the period 2002 to 2007 from three national circulation newspapers, Aftenposten, Dagbladet and VG, in print and on the

Table 6: Expectations regressions.

	<i>Non-fiction</i>		<i>Fiction</i>	
	Coeff	Std. Err.	Coeff	Std. Err.
lp_{11}	0.629	0.011	0.552	0.017
lp_{12}	0.153	0.012	0.147	0.018
lp_{13}	0.192	0.009	0.188	0.014
lp_{21}	0.226	0.011	0.286	0.016
lp_{22}	0.519	0.012	0.336	0.017
lp_{23}	0.230	0.009	0.269	0.013
lp_{31}	0.255	0.016	0.320	0.019
lp_{32}	0.290	0.017	0.164	0.020
lp_{33}	0.360	0.012	0.333	0.015
time ffx	yes		yes	
R^2	0.964		0.9565	
N	8730		4491	

Table 7: η correlations

	p_1	p_2	p_3
p_1	1		
p_2	0.43	1	
p_3	0.25	0.31	1

web. The reviews were accessed through Retrievers data base, a comprehensive Scandinavian media archive. About a third of the text reviews were editorially scored, the rest were unscored. Scored reviews are reviews that were given a score from 1 to 6 editorially by the newspaper. Just the fact that a title is reviewed is positively predictive of sales, yet there is additional information in the score itself. A very simple supervised machine learning algorithm was therefore applied to score also the unscored reviews. To get a feel for the impact of reviews on the demand patterns, I run some linear regressions of quantity on price with and without reviews. About 5% of the titles in the 2004 cohort are reviewed. I then run a

Table 8: Summary statistics

Variable	Mean	Std. Dev.
q	122.16	674.94
p	1.88	1.05
Reviewstate	0.05	0.21
N	39231	

simple linear regression of quantity on prices with and without reviews. Reviews are seen to be correlated with both prices and quantities. The bias from omitting review is towards zero as expected. The impact of reviews are about two thirds of the mean sales per observation, which seems economically significant. These regressions make no causal claims. Though reviews could certainly spur the sale of books, already well selling books may also be more likely to get reviewed. The role of the review variable in the analysis is simply to proxy for unobservable perceptions of quality that might affect both demand and retailer pricing.

In the columns to the right in (9), period and genre fixed effects are included. The period fixed effects has as expected a substantial impact on the price coefficients as it controls for the declining price path over the lifecycle.

Table 9: Reduced form text analysis with and without fixed effects


	Coeff.	Std. Err.	Coeff.	Std. Err.	Coeff.	Std. Err.	Coeff.	Std. Err.
constant	123.622	6.987	121.211	7.001	198.348	25.354	200.898	25.352
p	-0.774	3.243	-1.489	3.245	-18.523	3.748	-19.278	3.750
review	-	-	81.118	16.229	-	-	80.517	16.803
period ffx		no		no		yes		yes
genre ffx		no		no		yes		yes
N	39231							

C.1 Validation of the scoring algorithm

The algorithm was validated on a subsample of 82 titles which had at least one scored and one unscored review. Based on the results for the validation set, the algorithm display reasonable performance in adding information about the score. The odds ratio is $\frac{5}{4}$ for negative reviews and $\frac{4}{3}$ for the positive reviews. It is a limit to how well it can do with such a minimum of text that is furthermore typically complex. The score variable however picks up some variation beyond the binary information in being reviewed that has the expected sign and can proxy for unobserved demand shocks.

C.2 Reduced form evidence across genres and over years

Figure (10) shows that the price paths are fairly similar in all years following the deregulation. Prices start declining earlier than with RPM and fall to about the same level, which is indicative of a stable change in price strategies. Similarly,

**Figure 10:** Lifecycle price paths year by year.

the shift of demand from early at high prices to later at lower prices is also seen to be stable across years after the deregulation. For all the years following the deregulation, demand shifts towards lower future prices.

The price patterns are similar also across genres. Figures (12) and (13) show the price patterns broken down on Fiction and Non-Fiction, respectively. The price pattern is seen to be qualitatively similar.

D Demand elasticities and derivatives

The demand derivatives of the inside goods are

$$\frac{\partial D_{jt}}{\partial p_{kt}} = \begin{cases} \gamma_p \left(\sum_g w^g D_{jt}^g (1 - D_{jt}^g) \right) & \text{if } j = k, j = 1, \dots, J \\ -\gamma_p \left(\sum_g w^g D_{jt}^g D_{kt} \right) & \text{if } j \neq k, j, k = 1, \dots, J \end{cases}$$

The elasticities are $\eta_{jt} = \frac{\partial D_{jt}}{\partial p_{kt}} \frac{p_{kt}}{D_{jt}}$.




Figure 11: Lifecycle demand year by year.

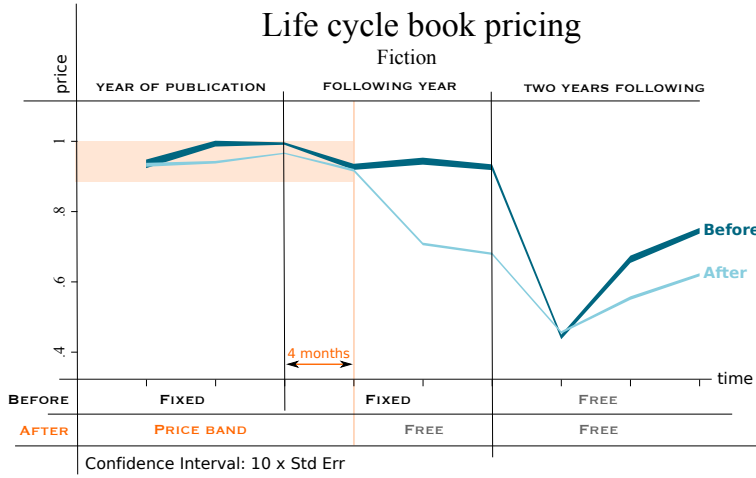


Figure 12: Lifecycle price paths Fiction.

The inter temporal elasticities are meant to measure the sensitivity of demand to a change in the expectation of future prices. It is not entirely clear what should be understood by a small change in expectations since the expectations are represented by a distribution F . It is taken to be a change in the first moment $\mu_{kt+1} = \int p_{kt+1} dF(p_{kt+1}|s_t)$. The inter temporal division ratios now depend on the derivatives of the option value wrt. expectations. These are used extensively in the counterfactuals. The expectations enter through the option value

$$v_{0t} = \beta \int \left(\ln \left(\sum_{l=0}^J \exp(v_{lt+1}) \right) + \Gamma \right) dF(s_{t+1}|s_t)$$

with derivatives $\frac{\partial v_{0t}}{\partial \mu_{k,t+1}}$ that are evaluated numerically using the chain rule and quadrature.

A special case of interest is perfect foresight when $\Sigma \rightarrow 0$. The derivative simplifies considerably so

$$\frac{\partial v_{0t}}{\partial \mu_{k,t+1}} = \beta D_{k,t+1} \frac{\partial v_{k,t+1}}{\partial p_{k,t+1}}$$

that is easy to evaluate. A sketch of a proof along with the closed form expressions are given below. An intermediate case where only $\sigma_k \rightarrow 0$ individually is

$$\lim_{\sigma_k \rightarrow 0} \frac{\partial v_{0t}(s_t)}{\partial \mu_{k,t+1}} = \int \frac{\partial g(s_1, \dots, s_{k-1}, \mu_k, s_{k+1,t+1}, \dots)}{\partial s_{k,t+1}} f(s_{-k,t+1}|s_t) ds_{-k,t+1}$$




Figure 13: Lifecycle price paths Non-Fiction.

which also simplifies calculation significantly.

D.1 Sketch of proof for the derivatives of the option value wrt expectations

The Markovian transition process $F(s_{t+1}|s_t) = MVN(\mu, \Sigma)$. We want to find the limit of the derivative of the continuation value v_{0t} with respect to $\mu_j, t + 1$ as $\sigma_j \rightarrow 0$. That is intuitively appealing, but a bit of work to establish. Set $j = 1$ for notational convenience. The continuation value is

$$\begin{aligned} v_{0t}(s_t) &= \beta \int \left(\ln \left(\sum_{l=0}^J \exp(v_{lt+1}(s_{t+1})) \right) + \Gamma \right) dF(s_{t+1}|s_t) \\ &= \int g(s_{t+1}) f(s_{t+1}|s_t) ds_{t+1} \end{aligned}$$

where $F(s_{t+1}|s_t)$ is $MVN(\mu, \Sigma)$. So the desired result is

$$\lim_{\sigma_1 \rightarrow 0} \frac{\partial v_{0t}(s_t)}{\partial \mu_{1,t+1}} = \int \frac{\partial g(\mu_1, s_{-1,t+1})}{\partial s_{1,t+1}} f(s_{-1,t+1}|s_t) ds_{-1,t+1}$$

To see that, suppose for now F is univariate $N(\mu, \sigma)$. Then

$$\frac{\partial}{\partial \mu_{t+1}} \int g(s_{t+1}) f(s_{t+1}|s_t) ds_{t+1} = \frac{\partial g(\mu_{t+1})}{\partial s_{t+1}}$$

Start by noting that from the symmetry of the normal $f'_\mu = -f'_s$. Take the derivative

$$- \lim_{\sigma \rightarrow 0} \frac{\partial v_{0t}(s_t)}{\partial \mu_{t+1}} = \lim_{\sigma \rightarrow 0} \int g(s) f'_s(s) ds + \lim_{\sigma_k \rightarrow 0} \int g'_s(s) f(s) ds$$

and use integration by parts in the first line

$$\begin{aligned} \lim_{\sigma \rightarrow 0} \int g(s) f'_s(s) ds + \lim_{\sigma_k \rightarrow 0} \int g'_s(s) f(s) ds &= \lim_{\sigma \rightarrow 0} g(s) f(s) \\ \int \lim_{\sigma \rightarrow 0} g(s) f(s) h(s) ds + \int \lim_{\sigma_k \rightarrow 0} g'_s(s) f(s) ds &= \lim_{\sigma \rightarrow 0} g(s) f(s) \\ \int \lim_{\sigma \rightarrow 0} g'_s(s) f(s) ds &= -g'_\mu(\mu) \end{aligned}$$

The second line uses that since f is normal, it can be written $f'(s) = f(s)h(s)$, where $h(s)$ is linear, and where the order of limits and integration on the left hand side is interchanged. By symmetry of the normal, $\lim_{s \rightarrow \infty} f(s) = \lim_{s \rightarrow -\infty} f(s) =$

0. The third line uses that f converges at an exponential rate, faster than g . Therefore $\lim_{s \rightarrow -\infty} g(s)f(s) = \lim_{s \rightarrow \infty} g(s)f(s) = 0$ and the right hand side is zero. Then $\lim_{\sigma \rightarrow 0} \int g(s)f(s)h(s)ds = \int \lim_{\sigma \rightarrow 0} g(s)f(s)h(s)ds = 0$, since the exponential rate of convergence of $g(s)f(s)$ dominates the linear divergence of $h(s)$. Finally, as $\sigma \rightarrow 0$, $f(s)$ becomes a Dirac delta spiking at μ . As $f'_\mu = -f'_s$, the result in the third line now follows.

The extension to multidimensional s is straightforward. In particular, when $\Sigma \rightarrow 0$ (perfect foresight over all prices), the partial derivative is just $\frac{v_{0t}}{\partial \mu_{kt+1}} = \frac{\partial g(\mu)}{\partial \mu_{kt+1}}$ with $\mu_{kt+1} = p_{kt+1}$, as expected. We then have that with perfect foresight

$$\begin{aligned} \frac{\partial D_{jt}^h(s_t)}{\partial \mu_{kt+1}} &= -D_{jt}^h D_{0t}^g \frac{\partial v_{0t}^h}{\partial \mu_{kt+1}} \\ &= -\beta \gamma_p D_{jt}^h D_{0t}^h \int \left(\ln \left(\sum_{k=0}^J \exp(v_{kt+1}^h(s_{t+1})) \right) + \Gamma \right) \frac{\partial f(s_{t+1}|s_t)}{\partial \mu_{kt+1}} ds_{t+1} \end{aligned}$$

for $j, k = 1, \dots, J$, and similarly for group l . The inter temporal elasticities at the market level are the weighted sum of each consumer groups elasticities

$$El_{\mu_{kt+1}} D_{jt}(s_t) = -\beta \gamma_p \mu_{kt+1} \sum_{g=h,l} \omega_{jt}^g D_{0t}^g \int \left(\ln \left(\sum_{k=0}^J \exp(v_{kt+1}^g(s_{t+1})) \right) + \Gamma \right) \frac{\partial f(s_{t+1}|s_t)}{\partial \mu_{kt+1}} ds_{t+1}$$

with weights $\omega_{jt}^g = \frac{D_{jt}^g}{\sum_g D_{jt}^g}$.