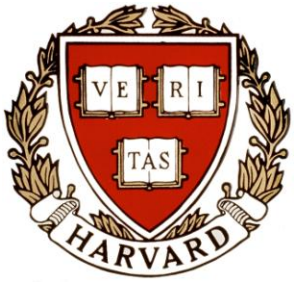

H I E R

Harvard Institute of Economic Research

Discussion Paper Number 2172

Welfare Effects of Full-line Forcing Contracts in
the Video Rental Industry

by

Justin Ho, Katherine Ho,
and Julie Holland Mortimer

May 2009

Harvard University
Cambridge, Massachusetts

Welfare Effects of Full-line Forcing Contracts in the Video Rental Industry*

Justin Ho†, Katherine Ho‡, and Julie Holland Mortimer§

December 17, 2008

Abstract

An extensive theoretical literature on bundling examines its potential efficiency and anti-competitive effects. In this paper, we provide an empirical study of bundling in a supply chain, referred to as full-line forcing. Using an extensive dataset on contracts between video retailers and movie distributors, we identify and measure three effects of full-line forcing: market coverage, leverage, and efficiency. We estimate a structural model of demand and the cost of holding inventory. Using the estimated parameters, we examine the effects of delaying the introduction of full-line forcing contracts. We find that delay results in a welfare loss for firms and consumers.

*We thank Lanier Benkard, Steve Berry, Richard Mortimer, Ariel Pakes, Mike Riordan, Mike Whinston and seminar participants at Carnegie Mellon University, Columbia University, the Cowles Summer Conference at Yale University, the Chicago-Northwestern IO and Marketing Conference, Duke University, Harvard University, MIT, New York University IO Day, Stanford GSB, Stanford Institute for Theoretical Economics Summer Workshop and the University of Wisconsin-Madison for helpful comments. We thank Yiwei Zhang for research assistance. The data for this study were generously provided to us by Rentrak Corporation. Any remaining errors are our own.

†Department of Economics, Harvard University, Cambridge, MA 02138. Email: jho@fas.harvard.edu

‡Department of Economics, Columbia University, New York, NY 10027 and NBER. Email: katherine.ho@columbia.edu

§Department of Economics, Harvard University, Cambridge, MA 02138 and NBER. Email: mortimer@fas.harvard.edu

1 Introduction

In vertically-separated markets, upstream firms face the challenge of inducing separately-owned downstream firms to carry their products and to resell them at prices that maximize profits of upstream producers. When managing these vertical relationships, upstream firms face both incentive-based participation constraints among downstream firms as well as legal constraints. Incentive-based participation constraints have led to a proliferation of vertical arrangements such as: resale price maintenance, exclusive dealing, exclusive territories, vertical contracts (e.g., revenue sharing), and bundling. An extensive theoretical literature has investigated the incentives for pursuing these vertical arrangements, and their potential impact on competition and welfare. These impacts have also made vertical arrangements an important focus of antitrust law. Despite this intense scrutiny, empirical evidence addressing the impact of different types of vertical arrangements is sparse. In this paper, we report empirical evidence on the welfare impacts of a revenue-sharing form of bundling in a vertically-separated supply setting. The bundling arrangement takes the form of a "full-line forcing" contract, in which firms are rewarded for accepting a producer's full line of products.

We consider the effects of full-line forcing or bundling contracts in the video rental industry. Movie distributors offer video rental stores the choice of three main contract types for distributing the bulk of their titles: linear pricing (LP), revenue sharing (RS), and full-line forcing (FLF). Under linear-pricing contracts, the store pays a fixed, upfront cost per tape, usually between \$65 and \$70. Under revenue-sharing contracts the upfront cost is much lower (around \$8-\$10 per tape) but the store also pays a fraction of the rental revenues (in the region of 55 percent) to the distributor. Full-line forcing contracts provide better revenue-sharing terms than the RS contracts (upfront costs of \$3 per tape and revenue-sharing payments of 35 - 40 percent), but require the store to buy minimum quantities of every title released by the distributor during the period of the contract (usually 12 months).[1]

We discuss three potential welfare effects of introducing FLF contracts in this application. First, a "market coverage effect" occurs if a store chooses a FLF contract when it would otherwise not have taken all of the distributor's titles, increasing the number of the distributor's titles that are available to consumers. Second, and conversely, a "leverage effect" occurs if the store compensates for the requirement to take all of a distributor's titles by dropping some titles produced by other distributors, particularly if inventory holding costs are high. Finally, an "efficiency effect" occurs if there are titles for which the store stocks higher levels of inventory in response to the lower upfront cost per tape on FLF terms compared to LP terms. Stores, on average, respond to the high upfront cost per tape under LP contracts by purchasing fewer tapes per title under LP than under RS. This may lead to queuing and other inefficiencies for LP compared to RS titles. However, RS mitigates the problem for low-value titles because stores tend to choose RS contracts for titles where they

[1] Distributors sell some additional titles on "sell-through pricing" contracts, in which all buyers, including video rental stores, can purchase tapes for around \$20-\$25 each. There is no contract choice for these titles, which usually include children's movies or titles with "teenager" appeal.

expect to have low demand. When FLF contracts are introduced, the store is required to take all of the distributor's titles on revenue-sharing terms, implying that some titles will be pulled out of LP contracts and into contracts with much lower costs per tape. This reduces the inefficiency from the store's low inventory choices under LP, and is the source of the efficiency effect.[2] The overall effect of bundling on efficiency and welfare depends on the relative importance of these three effects and is an empirical question.

We develop an empirical model of the industry and perform counterfactual analyses to investigate the three effects of bundling contracts. We ask how different the market would look in terms of the number of titles offered to consumers, the mix of distributors producing those titles, prices and store and distributor profits if the availability of FLF contracts was delayed. Our initial reduced-form analyses show that the overall patterns in the data are consistent with the market coverage and efficiency effects, but that the leverage effect may be small. Selection issues - in particular those caused by stores choosing which contracts to accept for different distributors - mean that a structural model is needed to analyze the market fully.

Using a new and extensive dataset, we model consumer demand for titles using a flexible nested-logit framework that takes advantage of the detailed transaction information in our dataset by including both store and title fixed effects, as well as decay rates, prices and numerous interaction terms. The demand system accounts for competition across titles, and allows the choice set for consumers to adjust in each month based on the set of new titles released by distributors. The missing link is then the cost of holding inventory: this is critical to our understanding of the market coverage and leverage effects. It includes rent, insurance and restocking costs, the value of reselling used tapes and also the potential value to the retailer of adding tapes in terms of drawing new customers into the store. We estimate this cost using a method of moments methodology with inequalities, following the approach developed in Pakes, Porter, Ho and Ishii (2007).

The counterfactual analyses examine the impact of delaying the introduction of the FLF contract of one of the participating distributors. The study indicates that average profits of participating stores fall slightly when FLF introduction is delayed; stores drop on average 1.8 of the five titles offered by this distributor and switch most of the other titles to RS contracts. The FLF distributor's profits fall substantially as a result. Competing distributors pick up a very small number of new stores. Overall, our results indicate that FLF contracts are, on average, welfare-improving, and that the leverage effect is outweighed by the market coverage and efficiency effects.

Understanding the effects of FLF contracts has very general applications, and speaks broadly to two literatures. First, the literature on bundling/tying, discussed below, focuses on the benefits that firms receive from these types of arrangements through their ability to mimic price discrimination or apply leverage across markets. Second, the literature on vertical arrangements focuses on the potential of such arrangements to both soften competition (through foreclosure or by raising rivals' costs) and induce efficient investments. All of these mechanisms may be present in FLF contracts, because they represent bundling in the context of a vertical arrangement. In the example we

[2]Note that this relates specifically to the revenue-sharing nature of the FLF contracts in this industry.

study, full-line forcing terms include a revenue-sharing component, making the vertical structure particularly salient. While these contracts may enhance overall profits for the upstream firms involved, understanding the welfare effects is an empirical question, because consumer welfare and the profits of downstream retailers can either increase or decrease.

The three effects that we discuss (market coverage, leverage, and efficiency) are impacted by both the tying and revenue-sharing components of the contracts we study. Tying affects market coverage (the decision to take a title), while revenue-sharing terms may further affect the number of tapes taken (i.e., the size of a store's inventory). Similarly, FLF contracts affect leverage, both through the decision to hold another distributor's titles as well as the inventory level of those titles. Finally, the efficiency effect depends on both the tying and the revenue-sharing aspects of the contracts as well: tying is required to induce firms to forgo less efficient LP contracts on high-value titles, and revenue-sharing terms are necessary to reduce upfront costs so that inventory levels are closer to what an integrated firm would choose.

Finally, the legal environment makes FLF contracts, and their welfare impacts, particularly interesting. As noted by Shy (1995), referring to U.S. law, "courts have been more receptive to vertical arrangements that did not involve price restraints." This is due in part to the potential for conflict between federal antitrust laws that govern price fixing and state laws that govern fair trade between firms. To the extent that firms can navigate their vertical relationships via non-price strategies (such as tying/bundling, revenue-sharing, quantity requirements, etc.), the potential for facing allegations of antitrust violations may be reduced. As a result, such strategies are widely-adopted in vertical settings in many industries, making our study an important first step for understanding the implications of these types of arrangements more generally.

One question in the theoretical literature on bundling addresses the reason why upstream firms might choose to offer bundling contracts. While our analysis does not focus on this particular question, the three effects discussed above (efficiency, leverage, and market coverage) provide potential explanations for distributors adopting bundling contracts. The theoretical literature does not consider the efficiency effect when considering the upstream firm's decision to offer bundling contracts, with the exception of Burstein (1960), which views full-line forcing as a means of achieving the effects of vertical integration. The leverage effect has been discussed in numerous theoretical papers including, for example, Whinston (1990), Choi and Stefanadis (2001), Carlton and Waldman (2002) and Nalebuff (2004). The market coverage effect is primarily a supply-chain effect, and is not addressed by the existing bundling literature, which focuses on bundling to end-consumers. Another potential explanation, and one that is addressed extensively by the theoretical literature, is price discrimination. If the preferences of the downstream firms for each good are negatively correlated, then an upstream firm can profit by bundling goods together. Adams and Yellen (1976) provide the first formal model of price discrimination through tying, building on a seminal paper by Stigler (1962). Their work was later generalized by McAfee, McMillian and Whinston (1989) and Salinger (1995) among others. Our demand model allows for negative correlation in preferences across stores when the consumer population differs across markets and preferences for particular

titles vary across demographic groups. However, our data, which cover seven distributors offering FLF contracts, do not allow us to analyze the price discrimination story explicitly. We focus instead on the store's choice of whether to take the bundling contract. The demand model is also rich enough to consider this question, since it allows the store to predict which titles will be particularly popular with its population.

Mortimer (2008) studied the impact of product-by-product vertical distribution arrangements on profits and consumer surplus, but did not address the issue of product bundling in the supply chain, which is the focus of this paper. The focus on mixed bundling requires us to extend the data used in Mortimer (2008) considerably–both extending the panel by two years in order to follow the adoption of the FLF contracts, and recovering detailed monthly activity for each title at each store. The bundling focus also requires us to implement a more flexible demand system that accommodates correlation in demand between titles, and allows for a more flexible cost structure at the retail level.

There are, to our knowledge, very few previous empirical papers on bundling. One paper, by Chu, Leslie and Sorensen (2007), studies bundling of tickets sold to consumers by a theatre company that produces a season of eight plays. The authors focus on examining the profitability of simple alternative pricing strategies to mixed bundling, and show that these alternatives can yield profits that are very close to those of mixed bundling. Crawford (2008) examines discriminatory incentives for bundling in cable TV, and Byzaalov (2008) and Crawford and Yurukoglu (2008) estimate the welfare effects of bundling in the retail market for cable television. There is also a small literature that uses reduced form analyses to investigate the pro-competitive and anti-competitive effects of slotting allowances (which are paid by manufacturers to supermarkets in order to reserve shelf space for their products). See for example Marx and Shaffer (2004). To our knowledge no prior papers have presented results for estimating a structural model of bundling in a supply chain setting, despite the importance of bundling features in many vertical arrangements. For example, FLF contracts (or related "softer" arrangements like rebates or slotting fees in exchange for carrying a manufacturer's full line) are common in many other industries, from groceries and vending to other retail markets, although many of these examples have not necessarily been characterized by revenue-sharing terms at the single product level.

This paper continues as follows. In Section 2 we outline the important institutional features of the industry and discuss the empirical implications of the theoretical literature on tying. Section 3 describes the data; Section 4 sets out our reduced form analysis. In Section 5 we provide an overview of the model. Section 6 considers demand, Section 7 covers the inequalities methodology and Section 8 describes our counterfactual analyses. Finally, Section 9 concludes.

2 Full-Line Forcing in the Video Rental Market

This section summarizes some important institutional features of the market and discusses the implications of tying for efficiency in this industry.

2.1 The video rental market

The video rental industry has two primary tiers: distributors, who distribute movies, and video rental stores, who acquire movies and offer them for rental and sale to consumers.[3] Three different contractual forms are used to distribute titles targeted to the rental market from distributors to rental stores. The first is linear pricing. Under linear-pricing contracts, a store purchases a title from the distributor for a fixed cost per tape, usually between $65 and $70. They may also offer quantity discounts (introducing some second-degree price discrimination).

The second contractual form is revenue sharing. Under revenue-sharing contracts a rental store leases a title from a distributor for a low upfront cost per tape, but receives a share of the revenues generated by a title. In the typical RS contract, the distributor charges an upfront cost of around $8 per tape and receives about 55% of the rental revenue. The inventory decision of the rental store is often constrained by both maximum and (often binding) minimum quantity restrictions. RS and LP contracts are both implemented on a per-title basis. That is, for each individual title, the rental store chooses both whether to purchase the title and the contractual form.

In contrast, the third form of contract, the full-line forcing contract or output program, requires the rental store to purchase all titles released by the distributor during the period of the agreement (typically 12 months) and to take them all under the same contract type.[4] In many other respects FLF contracts resemble RS contracts. For each title taken on full-line forcing terms, the distributor receives an upfront payment per tape and a share of the revenues, both of which are usually lower than under RS contracts. The quantity taken by the retailer is again restricted to be within a range, where the lower bound (on average eleven tapes per title) is again frequently binding.

Distributors have only a limited ability to price discriminate across retailers. The 1936 Robinson-Patman Act prevents distributors from offering different prices to competing buyers for exactly the same product. Furthermore, copyright laws permit stores to freely resell tapes purchased from distributors under LP contracts, which effectively limits their ability to use second-degree price discrimination. We do observe a few volume discounts under LP contracts and some negotiated deals under RS (again based largely on volume), and our estimation takes these into account. There may be some volume discounts for LP contracts that we do not observe; we include the effects of these discounts on stores costs in our inequalities analysis. The maximum and minimum quantity requirements for RS contracts also vary by distributor with the box office of the movie and the size of the store: this variation is observed in our data and accounted for in our model.

In addition to setting terms for each contractual form, the distributor can, in theory, choose which contractual forms to offer. In particular, one might expect the distributor to choose not to

[3] Rentrak receives a small cut of the distributors' profits under RS and FLF contracts. While it may have a role in trying to persuade distributors to offer RS and FLF contracts, it does not influence stores' choices conditional on the contract types offered for each title. We therefore exclude Rentrak and wholesalers from our model.

[4] Some exceptions apply: titles released by the distributor on sell-through pricing contracts are exempt, and several distributors allow for limitations on the total number of titles that a retailer must accept within any given month. Usually, this limit is three titles per month: if the distributor releases more than three titles in a month (a rare event), the retailer is only obligated to accept three of them. Finally, FLF contracts also typically include opt-out clauses for movies with 'objectionable' content.

offer LP contracts since these are the least flexible contractual arrangements. In reality, however, RS contracts were not widely used before the end of 1997, and FLF contracts were not introduced until the middle of our dataset, in February 1999. One likely reason is that both contractual forms require extensive computer monitoring of millions of transactions, and only about half of the stores in the industry had the technology to adopt these contracts by 1998.[5] Thus, eliminating LP contracts during the period of our data may have substantially reduced distributors' target market.[6] This implies that rental stores can discipline the distributors by opting to take LP contracts when revenue-sharing splits are not satisfactory. The empirical evidence suggests that LP contracts continue to be offered to all firms even when RS contracts are also available.

One further institutional detail concerns "sell-through priced" titles. These include, for example, children's movies and some very popular titles: the distributor sells these movies to all buyers, including video rental stores, for relatively low prices, often around $20-25 per tape.[7] There is no contract choice for sell-through priced titles: we condition on these titles' existence in the demand model and account for them in our calculation of the store's total returns in the inequalities framework but we do not model the distributors' contract choice.

Finally, the sales market is important for distributors and should be included in any model of their choices of contract types. However, sales provide only a small proportion of total revenues of rental stores, whose choices are the focus of this paper.

2.2 Empirical Implications of Theory on Tying

We consider three potential welfare effects of introducing FLF contracts. First they affect retailers' inventory choices. The high cost of tapes under LP contracts causes stores to choose low inventory levels for LP titles compared to the inventory choice of a vertically-integrated firm. This inefficiency is reduced when titles shift to RS contracts because the average upfront cost per tape falls and the store's inventory level increases.[8] However, when only RS and LP contracts are available and demand is independent across titles, retailers will choose LP contracts when expected demand for the title is relatively high.[9] Thus, offering RS contracts in addition to LP contracts may not mitigate the efficiency loss from low inventory choices for high-value titles, for which the loss is relatively large. This is the source of the efficiency effect of a FLF contract: since the contract requires the store to take all of the relevant distributor's titles under terms that include a low

[5] Our dataset includes only stores that have the technology to do revenue-sharing term.

[6] Distributors are effectively prevented from offering LP contracts only to stores without revenue-sharing capability by the Copyright Act of 1976. This states that the owner of a lawful copy can "sell or otherwise dispose of" the copy and implies that retailers with the ability to participate in revenue-sharing agreements cannot be excluded from choosing LP contracts for a title unless LP contracts are not offered to any retailers for that title.

[7] Sell-through priced titles are exempt from the requirement that stores choosing a FLF contract take all of the distributor's titles on FLF terms.

[8] In a regression of per-tape retailer payments to the distributor on contract type, title fixed effects and store demographics, the average payment per tape taken on a RS contract was $34 lower than the average under LP contracts (standard error 0.02). The average under FLF contracts was $41 lower than that under LP contracts (standard error 0.06).

[9] Mortimer (2007) demonstrates this in a market that is consistent with the assumptions in our empirical model.

upfront cost and a low average cost per tape, valuable titles are pulled out of LP contracts, and this may significantly reduce the low inventory problem.

The low inventory effect of LP may not substantially affect the prices charged to consumers because there are two opposing effects. First, the selection of high demand titles under LP contracts (and possibly the higher average cost per tape) provides an incentive to increase the average rental price compared to RS or FLF titles. However, once the inventory has been purchased, the store has an incentive to price LP titles below the RS or FLF titles that compete with them, in order to draw consumers to the titles for which they capture 100% of the rental revenues. These two offsetting effects may imply small differences in rental prices between contract types.

There are two other potential welfare effects of introducing FLF contracts. First, if the store previously took only a subset of the distributor's titles, the fact that it must now take all of them implies a positive effect on market coverage. This is probably welfare-improving since it increases the size of consumers' choice sets. Conversely, this effect together with the non-zero cost of holding inventory may prompt the store to drop other distributors' titles: this is the leverage effect and is likely welfare-reducing since it reduces competition between distributors.

The relative magnitudes of these three effects will depend on the mean and variance of demand for the titles released by different distributors and the extent of complementarities between them and also on stores' inventory holding costs. The aggregate welfare effect of FLF contracts is therefore an empirical question.

3 The Dataset

Our primary data source is Rentrak Corporation, an organization that distributes movies under revenue-sharing and full-line forcing contractual arrangements. They also monitor all titles at participating retailers, and facilitate payments between retailers and distributors under the RS and FLF contracts. The complete dataset combines information from previous studies (Mortimer 2007 and 2008) with additional information from Rentrak on FLF contracts. Over 11,000 retailers used Rentrak between 1998 and 2001, accounting for over half of all retailers in the industry. Approximately 4,000 of these are Blockbuster Video and Hollywood Video stores: we do not observe their transactions. We observe 7,525 retailers (over 30% of all stores in the industry), ranging in size from single-store locations to a chain with 1,652 locations. For each store we observe transaction data between January 1, 1998 and June 30, 2002 and follow 1,025 titles released during this period.[10]

For each store we observe the total monthly revenue of a store, its zip code, the size of its chain and considerable detail regarding product mix, such as the overall percentages of game, adult, rental, and sales revenues. We also observe the date the store joined the Rentrak database and the date the store left Rentrak if applicable. The vast majority of store exits (over 90 per cent) represent

[10]Full-line forcing arrangements continue to be an important contractual form in the industry for the DVD format. However, unlike for VHS in the period that we study, the focus of FLF contracts for DVDs has been older blockbuster titles or "classic" movies.

store closure.[11] The zip code information allows us to supplement the primary Rentrak data with several additional sources. Phonebook listings of competing video retail locations in each year, as well as separate indicators of competing Blockbuster and Hollywood Video locations are included. We also merge in data from the 2000 US Census on the local demographic characteristics of each store. We define a local market as a zip code area: the average zip code contains approximately 24,000 people and 2.6 video retail stores. Larger areas, such as 4-digit zip codes or Metropolitan Statistical Areas (MSA's) are also feasible ways of defining markets but are probably too large for most video store customers.

Every movie title is tracked individually, using a title identifier but not the actual title name. For each title we observe a distributor identifier (but not the actual distributor name), its month of release to video, genre (i.e., Action/Adventure, Children/Family, Comedy, Drama, Horror, Romance, and Sci-Fi), and MPAA rating (G, PG, PG-13, and R). We also observe box-office categories, denoted A, B, C and D. Titles in the A category have theatrical box-office revenues of more than $40 million; those in the B and C categories have revenues of $15-40 million and $0.25-15 million respectively. Titles in the D category do not have a theatrical release: these are "direct-to-video" titles such as instructional or exercise videos. Many of these titles are bought only by a single store; we exclude D titles from our analysis. The dataset includes 212 A titles, 195 B titles, and 618 C titles.

In addition to title characteristics, we observe the terms of the RS and FLF contracts offered to retailers for each title, as well as retail prices under LP and sell-through pricing contracts. Rentrak does not provide the actual wholesale prices paid by retailers under linear-pricing terms: we adjust the retail price to reflect the true wholesale price using guidance from Rentrak and industry sources (see Mortimer 2007 for details).

Finally, at the store-title level we observe the type of contract chosen by the retailer and the number of tapes purchased. Transaction data are recorded at the store-title-week level, and provide information on the number of rentals per tape, total weekly revenues per tape, and inventory levels (which do not vary across weeks). We discard observations for titles released after January 1, 2002 so that rental activity for each title is tracked for at least 6 months. We aggregate weekly rental data to the month level (both the number of rentals and average rental prices for the month) in order to smooth out any weekly demand fluctuations. We therefore have 54 months of transaction data for titles released over 48 months.

We take several steps to clean the dataset. First we exclude observations where average price per rental is less than $0.50 or more than $7 and those where store demographic data are missing. We drop five titles whose wholesale price is zero. Ten titles have two values for release month: for nine of them the majority of observations have the same (earlier) value so we assume that the later date refers to a special edition and switch to the earlier date for all observations. The tenth title has half the observations with one release date and half with another; we drop this title from the

[11] For 1,116 stores, data collection ended for titles released after December 1999. We include these stores in both our demand and supply models up to that date.

dataset. We are left with 6,393 stores, 961 titles (201 in the A box-office category, 188 B titles and 572 C titles) and 59 distributors in the dataset.

FLF contracts first appear in the dataset in February 1999. Seven out of 59 distributors offer a FLF contract at some point in our panel. The average number of titles released per year by these seven "FLF distributors" was 10.5; the maximum was 24 and the minimum 1. The 52 "non-FLF distributors" released on average 7.4 titles per year, with a maximum of 46 and a minimum of 1. Of the 6,393 retail stores in the clean dataset, 6,356 participated in at least one LP contract during the period of the analysis, 6,120 participated in at least one RS contract, 5,077 participated in at least one FLF contract and 6,171 participated in at least one sell-through pricing contract. On average stores take 46% of the titles released per distributor-year. The proportion is higher for FLF distributors (59%, or 53% if we include only stores that have no FLF contracts with the relevant distributor) than for other distributors (43%). That is, FLF distributors both release more titles per year and also have higher take-up rates than other distributors.

Stores are categorized into ten sizes, called "tiers," with tier 1 containing the smallest stores and tier 10 the largest stores. The average number of titles taken per month increases with store tier from 9.6 in tier 1 to 12.9 in tier 6. It then falls in each subsequent tier to a low of 9.3 in tier 9, and back up to 9.5 in tier 10. The average number of tapes taken per month increases with every store tier, from 43.0 in tier 1 to 307.6 in tier 10. "FLF stores" take an average of 1.6 more titles per month (11.9 vs. 10.3), and 47.6 more tapes per month (132.2 vs. 84.6) than "non-FLF stores."

Additional summary statistics are provided in Tables 1 to 4. Table 1 sets out average contract terms, numbers of rentals, prices and inventories for each contract type. Averages are taken across store-title pairs. The average estimated wholesale cost for LP contracts is \$66.82, compared to an average upfront cost of \$8.47 for RS contracts, \$3.62 for FLF contracts and a cost of \$15.17 for sell-through priced contracts. Retailers on average keep 46% of revenues under RS contracts, and 59% of revenues under FLF contracts. The minimum number of tapes per title is 10 on average for RS contracts and 11 for FLF contracts. On average, the maximum number of tapes allowed per title is 23 for both contracts. Average month 1 rentals are highest under RS contracts but the decay rate is also greatest for these titles; by month 3 LP titles have higher demand and this remains true in months 4 and 5.

Average inventory levels are highest for titles purchased under sell-through pricing and RS contracts and lowest for those under LP contracts. This is the source of the efficiency effect described above. Not surprisingly, retailers also extract the largest number of rentals per tape for titles purchased under LP contracts. As discussed in Section 2, the effect on average rental prices of the lower cost per tape under RS contracts compared to LP contracts is offset in part by the effect of the revenue-sharing component. However, a regression of price on indicators for months since release to video, by contract type, indicates that prices fall faster for RS titles than for LP titles. This may indicate that the price-increasing effect of a high cost per tape under LP slightly outweighs the opposing effect of the two-part tariff under RS. It also implies that the margin on which prices adjust may be the timing of removal of the "new release" sticker, with concurrent

price reduction or increase in the rental period (and resulting decrease in late fees collected).

Tables 2 and 3 summarize the number of titles released by distributors, and taken by stores, under different contract types. The majority of titles in our data were offered under LP contracts; approximately 56% were also offered under RS contracts. No FLF contracts were offered in the first year of our data; a total of ten were offered in year 2, eighteen in year 3, and 38 in year 4. Table 3 further reports that stores on average took many more titles on LP contracts than on other contract types.

Finally, Table 4 provides information on the size distribution of stores choosing different types of contracts. We begin by calculating the percent of each store's titles that were taken under each contract type. We then break down this distribution into quintiles and report, in the first panel of the table, the average store size (tier) for each quintile. The results demonstrate that stores that accept very few titles on LP contracts (the lowest quintile) are the small stores - these stores take a relatively high proportion of their titles on RS contracts. The stores that accept a high proportion of their titles on LP contracts are on average larger stores. This is consistent with the adverse selection effect noted above: large stores tend to be located in high-demand markets and therefore expect high demand for their titles. LP contracts are most profitable for these stores. The pattern for FLF contracts is similar to that for LP: larger stores are more likely to accept a high proportion of their titles on FLF contracts.

The second panel of the table looks at these patterns in more detail. For each quintile and contract type, we examine the percent of stores in that quintile/contract type that are store tiers 1-3 or store tiers 7-10. We normalize these percentages by the overall percent of stores that are in those tier groups across all quintiles and contract types. Thus, the result of 1.59 for tier 1-3 under LP and Quintile 1 indicates that store tiers 1-3 (small stores) are relatively over-represented in the first quintile of LP contracts (a value greater than 1 indicates over-representation, and a value less than 1 indicates under-representation). Overall, small stores are over-represented in the first, second, and third quintiles of LP contracts, the first and second quintile of FLF contracts and the third, fourth and fifth quintiles of RS contracts. The reverse pattern holds for large stores: these are over-represented in the fourth and fifth quintiles of LP and FLF contracts and in the first and second quintiles of RS contracts. However, similar to small stores, large stores are also over-represented in the first quintile of FLF contracts.

4 Reduced Form Evidence

We now discuss preliminary evidence and patterns from the data. In particular, we ask whether reduced-form analyses can provide any evidence on the importance of the efficiency, market coverage and leverage effects of full-line forcing contracts.

4.1 Retailer Performance Across Contract Type

First we test the prediction that retailers who expect a relatively low draw of demand for a particular title will choose a revenue-sharing contract while retailers who expect high demand for that title will choose linear pricing. Demand for titles taken on full-line forcing contracts is likely to be somewhere in between that for revenue-sharing and linear-pricing contracts. The summary statistics above indicate that large stores (which tend to have high demand) are most likely to choose LP contracts, small stores are more likely to choose RS contracts and the stores choosing FLF contracts are similar to those choosing LP. We expect a similar pattern here except that FLF titles may generate lower revenues than LP titles since stores cannot "cherry pick" titles under FLF. We regress revenues for a store-title pair on an indicator for the adoption of a RS contract, an indicator for the adoption of a FLF contract and title fixed-effects. Consistent with our prediction, we find that revenues for Box Office category A titles are approximately \$273 lower under RS than under LP contracts (standard error of 4.18) and that revenues under FLF are \$189 lower than those under LP (standard error 16.30).[12] The relationship between RS and LP revenues is similar for Box Office category B and C titles.

Next we investigate in more detail which types of stores choose FLF contracts. A logit regression of a dummy for participation in these contracts on observable store characteristics indicates that larger stores and those in rural and suburban areas are more likely to adopt FLF contracts. Stores in markets with a high percent black or percent female population are less likely to adopt FLF.

Finally, we would like to investigate whether the introduction of RS and FLF contracts had a positive effect on market coverage. First the statistic noted earlier, that stores on average take 53% of titles released by distributors that offer FLF contracts at some point, excluding FLF contracts themselves, is consistent with a potentially large effect. We also regress the number of titles taken per distributor-month on an indicator for active FLF contracts for the relevant store-distributor pair and store, distributor and month fixed effects at the store-distributor-month level. We include only store-distributor pairs for which a FLF contract exists at some point in our panel and only months when the distributor offers FLF contracts. We are therefore looking within-store and asking whether taking a FLF contract from a particular distributor is correlated with title choices specific to that distributor, ignoring average distributor and month effects. The coefficient on FLF activity in this regression is positive and significant (coefficient 0.37, standard error 0.007). We then run similar regressions at the store-title level where the dependent variables are the number of tapes per title and the number of transactions per title. Here we add title fixed effects to the list of explanatory variables. The coefficient on FLF activity in the regression considering the number of tapes per title is 2.30 (standard error 0.13). This implies a positive market expansion effect in terms of both the number of titles taken and the number of tapes per title. The equivalent coefficient in the transactions regression is also positive and significant (coefficient 6.87, standard

[12]The results imply that the effect of low inventory choices for LP contracts, which are likely to reduce demand for these titles, is outweighed by the quality difference generated by the difference in tariffs between LP and RS contracts.

error 1.59). The new titles taken under the FLF contract generate more rentals than other titles, perhaps indicating that titles taken under FLF contracts have fewer stockouts than other titles on average.

4.2 Full-line Forcing and Competing Products

Our final reduced-form analysis investigates the leverage theory: that full-line forcing can have anticompetitive effects in the upstream market by reducing retailers' orders from other distributors. We might expect this effect to generate a negative correlation between the adoption of FLF contracts by a retailer and the orders (or rentals) of products from other, non-bundling distributors. However, most of the theories that generate such predictions consider full bundling rather than mixed bundling. In our application large stores' selection into different contractual forms in particular months and for particular distributors or titles may alter the intuition. In addition, these selection effects potentially bias the results of the reduced form analysis so a more detailed model is needed to estimate the magnitude of the leverage effect.

First we regress the number of titles taken per distributor-month on an indicator for active FLF contracts in that month with some other distributor at the store-distributor-month level. We include store, distributor and month fixed effects. For each store, we exclude from the regression distributors with which the store ever has a FLF contract and all months before FLF was offered by any distributor. The coefficient on the number of titles taken per distributor-month is positive and significant (coefficient of 0.098, standard error 0.001). As above we then run similar regressions at the store-title level where the dependent variables are again the number of tapes per title and the number of transactions per title, adding title fixed effects to the independent variables. The coefficient on the FLF dummy in the specification using the number of tapes per title as the left-hand side variable is 0.32 (standard error 0.05); that on transactions per title regression is 7.38 (standard error 0.76). These very simple regressions indicate that for a given store, the total number of titles taken per month from non-FLF distributors in months when the store takes FLF from some other distributor, excluding average month and title effects, is higher than the number taken in months when the store does no FLF. This result holds for both the number of tapes per title and the number of transactions per title. That is, the leverage effect seems to go in the opposite direction from that predicted by the theory.

These results likely reflect selection effects. While we include store fixed effects to address bias due to retailers selecting into FLF contracts, and distributor and month (and where possible also title) fixed effects to absorb the average effect of the numbers and qualities of titles released in particular months or by particular distributors, it is not feasible to include interactions between these fixed effects. The results may be consistent with stores taking FLF in periods when there is a shock to demand. For example, periods in which there is entry or exit of a close competitor. The structural model, described below, is needed to separate this and similar demand effects from the leverage effect we wish to identify.

5 Overview of the Structural Model

The summary statistics and reduced-form analyses provided some evidence that FLF contracts may affect the efficiency of contract types chosen for particular titles. The analysis is also consistent with the market coverage effect and indicates that the leverage effect may be small or even negative. A structural model is needed to correct for the selection problem arising from stores choosing which contract types to take in particular months. As noted, this selection issue makes it difficult to predict the effect of introducing new FLF contracts from the reduced form results alone. A structural model also allows for an analysis of overall welfare effects of FLF contracts.

The modeling approach we use has three elements. First, we estimate a demand system that parsimoniously captures the demand interactions between titles and across title categories. Since the impact of FLF contracts is to change the composition of the choice set, the inventory of each title and the price per rental, the focus of this demand system will be to capture the impact of adding or removing a title from the consumer's choice set and changing inventory and price. The second step is to use moment inequalities to infer the store's cost of holding inventory. The third step is to use the estimated model to run counterfactual experiments to infer the impact of FLF contracts on the division of the surplus between distributors and stores and on consumer surplus.

Before discussing the details of the model it is helpful to consider the factors the retailer takes into account when choosing between contract types. First, and most obviously, the price and revenue-sharing terms of the contract are important. As noted above, LP contracts are likely to be chosen in preference to RS when the store expects high demand for a title. Second, the retailer considers the inventory restrictions for both RS and FLF contracts: in both cases the minimum quantity restriction is on average higher than the average number of tapes per title taken under LP contracts and is often binding in our data. The restriction affects retailers both by increasing the cost of taking the title (the number of tapes to be purchased and the cost of storing tapes) and also potentially by increasing the expected level of demand for the title. For example, a higher inventory level implies a higher number of tapes on the shelf and therefore a title that is more visible to consumers. It may also act as a signal of high quality or a blockbuster title. Thus, high inventory may lead to high initial demand: consumers find out about and choose to rent the title more quickly than they would have done otherwise. This may also affect later demand for the title because of a durable goods issue: if a consumer rents a title in one month he is unlikely to rent it again later. We allow for this by including both inventory and inventory-month interactions in our demand model.[13]

There are three other potential effects of an increase in inventory. First, increasing the number of tapes taken per title (or the number of titles) may increase retailer profits by attracting new consumers to the store, inducing them either to switch from other video rental stores or to enter the market for the first time. Second, if the inventory of title X is increased and some other title Y

[13] A second mechanism through which inventory affects demand is stockouts. We have no data on stockouts and therefore cannot fully model this issue. We estimate demand at the store-title-month level, therefore allowing consumers to substitute intertemporally within a month when a stockout occurs.

is a close substitute released in the same month, then its rentals are likely to fall which could imply an overall reduction in retailer profits for this pair of titles.[14] Finally, if consumer preferences are correlated across months, then a change in X's inventory level in month 1 may affect title Y rentals in later months and this too may impact retailer profits.[15] Our model captures the second of these three effects; the impact of the other two effects on retailer profits is included in the cost of taking an extra tape, estimated in the inequalities analysis.

It is worth noting here that stores do not in general face physical inventory constraints. If they come close to running out of shelf space when storing tapes title-page-forwards, they simply store them spine-forwards (starting with the oldest titles). If they run out of space again they can hold some tapes under the counter or in a back room. Thus, the choice of inventory levels affects demand but not the number of titles that can be displayed.

6 A Model of Demand

6.1 Demand Methodology

The data provide information on the number of rentals and the total revenues for each title-store in each week. We aggregate this weekly information to the month level for two reasons. First, stockouts can lead to shifts in observed weekly transactions that are unrelated to true demand; allowing consumers to substitute across weeks within each month mitigates this problem. We do not observe periods of stockouts, which can include tapes that are lost or returned late. This is primarily a limitation that results from the rental nature of the product. Thus, it is difficult to implement the corrected demand estimator proposed in Conlon and Mortimer (2008) to account for stockouts explicitly. Second, we account for the changing set of competing titles due to the release of new titles over time: this would not be feasible in a weekly framework.[16] Our methodology for aggregating to the month level is as follows. For any title released in month one, we summarize over weeks 1 through 4 to generate month 1 demand.[17] Similarly, we summarize over weeks 5 - 8 to generate month 2 demand, weeks 9 - 13 to generate month 3 demand, and weeks 14 - 17 for demand in month 4. Finally we aggregate all remaining weeks into a "months 5 and above" observation. Approximately 84% of all rentals occur in the first 4 months after a title's release to video.[18] We construct prices at the monthly level by dividing monthly revenues by monthly

[14] For example, this would be the case if title Y was taken on LP or sell-through pricing terms but X was taken on RS or FLF so that the proportion of the rental revenues captured by the store was higher for Y than for X.

[15] Again the impact of this substitution effect on store profits depends on the contract types of titles X and Y.

[16] Using weekly demand when the choice set changes every week would require that we estimate over 35 different choice sets for each title if we allowed for four months of activity per title. This is computationally impossible for titles that are held by a small number of retailers, and difficult even for widely-held titles, when we wish to incorporate sensible decay patterns.

[17] We replace missing or negative values for weekly revenues or transactions with zeros.

[18] Note that titles released in the last week of a month will be tracked for just one week in the first month rather than 4 weeks, which will bias down the demand estimates for those titles. Fortunately, there appears to be no correlation between this and distributors, genres, or any other observable characteristic of movies, and based on industry discussions, we assume that this form of truncation is random. In addition, the "month 5+" revenues for titles released later in our time period will be smaller than those for titles released earlier.

transactions. Finally, we drop store-title-month observations with zero rentals/revenues.

We define a title's competitors in each month as the titles that were released during the previous 4 months (including the current month). This implies an assumption that titles released more than 4 months ago do not substitute for current releases. Only titles in this moving window are included in the analysis for the relevant month.[19,20]

We estimate a nested logit model of demand with nests defined as genre/box office class groups (i.e., "A" comedies).[21] The demand equation is:

$$u_{ijmt} = \delta_{jmt} + \zeta_{igmt} + (1-\sigma)\varepsilon_{ijmt} \tag{1}$$

where i indexes consumers, j titles, m stores, t months and g the genre/class group of the title. The term ζ_{igmt} is an idiosyncratic preference term common to all titles in group g and ε_{ijmt} is an idiosyncratic preference term specific to consumer i and the product indexed by jmt. Cardell (1997) gives conditions such that $[\zeta_{igmt} + (1-\sigma)\varepsilon_{ijmt}]$ has an extreme value distribution with $\sigma \in [0,1]$ parameterizing the correlation of the idiosyncratic preferences within group ($\sigma = 0$ means no correlation; $\sigma = 1$ means perfect correlation). Price varies across titles, geographic markets and months. The term δ_{jmt} is specified as:

$$\delta_{jmt} = \delta_j + \gamma_j z_m + \eta_m + \theta_t + \beta_t x_j + \lambda_t c_{jm} - \alpha p_{jmt} + \xi_{jmt} \tag{2}$$

where δ_j is a title fixed effect, η_m is a store fixed effect, θ_t is a month fixed effect, p_{jmt} is the average price per rental of the tape at store m in month t, and c_{jm} is the inventory of title j at store m. The last term ξ_{jmt} captures any unobservable quality of renting title j in market m in month t. This could include things such as local promotions of a particular movie in a month. We interact title dummies with store characteristics: these describe the demographics of the store's market. The variables are the percent white, the percent single and the percent with children. We therefore permit each store to predict the demand for a particular title based on the demographics of local consumers.[22]

[19] Our full dataset includes titles released between months 1 and 54 of our panel. The vast majority of rentals occur in the first 6 months after a title's release. However, for titles released after month 48, we do not observe 6 months of rental activity, so we restrict attention to titles released between months 1 and 48. Furthermore, for a store present in all 48 months, titles released between months 1-4 and 45-48 compete with those released in months -3 to 0 and 49-52 respectively, which we do not fully observe. We therefore exclude months 1-4 from the final demand and inequalities analysis in addition to titles released after month 44, ensuring that we include only months for which we observe the full choice set and at least 5 months of rental activity. Similarly for stores that enter or exit the sample, we exclude the first and last 4 months.

[20] We choose to pool the data across months rather than estimating demand separately in each month because the variation in choice sets offered across months enables us to identify a detailed set of interactions with the decay rate; see below. This approach also requires fewer normalizations: if we estimated month-by-month, each month would have an outside good which would need to be normalized to zero to enable cross-month comparisons.

[21] This can be interpreted as a random coefficients model with the random coefficients on group dummies. See Berry (1994) for a discussion. Since genre and box office category are the main sources of differentiation between titles (and 2 of our 3 observable title characteristics) this seems a natural structure for modeling heterogeneity in consumer preferences.

[22] We could also have interacted store dummies with title characteristics. We choose not to do this partly because our title characteristics are not very informative - see below for a discussion. In addition, the implied effect, that the

The decay rate θ_t captures two effects. The first is the simple idea that demand for a title falls over time as advertising and word-of-mouth "buzz" decrease. The second is the durable goods issue noted above: if a consumer rents a particular title in month 1 he is unlikely to be in the market for the same title in month 2.[23] We would ideally account for this effect by including title-month fixed effects, allowing for a completely flexible decay rate for each title; unfortunately the number of titles is too large for this to be feasible. Instead, we interact month fixed effects with title characteristics (box office class, genre, and rating) and double and triple interactions of these three groups of variables).[24] This implies constraining the decay rate to be the same for all titles in a particular box office category-genre-rating cell.[25] Finally, we also interact the decay rate with the store's inventory level for the particular title. This accounts for the different average inventory levels associated with different contract types.

Integrating out the idiosyncratic preference terms yields the following equation for estimation:

$$\ln(s_{jmt}) - \ln(s_{0mt}) = \delta_j + \gamma_j z_m + \eta_m + \theta_t + \beta_t x_j + \lambda_t c_{jm} - \alpha p_{jmt} + \sigma \ln(s_{jmt/gmt}) + \xi_{jmt} \quad (3)$$

where $s_{jmt/gmt}$ is the share of title j within group g at store m in month t. The outside option (with share s_{0mt}) is doing something other than watching a new release movie. Its share is calculated from a market share assumption: we assume that the market size (denoted M) is equal to 4 movie rentals per month per household in the store's zip code.

One further aspect of the data complicates the estimation process: we very rarely observe more than one store per zip code (although we do know the number of stores that exist in each zip code). We therefore cannot explicitly include the whole choice set in the demand estimation. We address this by treating each store as a monopolist in its market. If N stores actually exist in the market (according to the phone book) we assign $\frac{1}{N}$ of the total population to the observed store; we model demand for the store as coming from just that subset of consumers. This implies an assumption that stores in the same market are identical and have independent populations of potential customers; a change in characteristics might attract more customers from that population but would not steal business from other stores. We interpret this as an assumption that, when a consumer visits a video rental store, if he does not find the title he is looking for he will either rent something else or go home rather than visiting a different store. The relevant dimension of competition, particularly since we are considering bundling, is that across distributors within a store rather than that across stores. We model the former carefully but do not go into details on the latter.[26] In reality, if one

"quality of a store" differs across types of movies - would identify essentially the same effect as the $\gamma_j z_m$ term: that stores serving different demographic groups expect different movies to be popular.

[23] There is also a potential seasonality effect: a title released in December may have different demand from one released in June. This effect is absorbed into the title fixed effects, although not differentially across months.

[24] We include only interactions for which there are non-trivial numbers of observations. For example, there is only one PG action/adventure rated movie so we combine that cell with PG13 action/adventure movies.

[25] There is one further issue which we would ideally account for by including title-month fixed effects. If title A was introduced in month 4, it competed with and therefore affected demand for title B in month 7. It therefore had an impact on residual demand for title B in months 8-10. These interactions between months would be perfectly accounted for if we had a fully flexible time trend for each title.

[26] We could alternatively have included all observed stores in each market in the estimated demand system and

store improves its offering over time by adding titles or tapes, it may gain market share from other local stores. This effect will be identified in the inequalities analysis. However, we do not model other aspects of inter-store competition such as pricing and specific portfolio choices. One obvious concern is with Blockbuster, which has FLF contracts for a large number of titles and frequently has a larger portfolio than its competitors. We treat Blockbuster like any other store in the demand equation (in that, if there are 2 non-Blockbuster and 1 Blockbuster stores in the market, then each observed store's demand is predicted assuming a population $\frac{1}{3}$ of the total in the market). The store fixed effects absorb any differential effect that a Blockbuster dummy would have on demand, absent entry or exit by Blockbuster outlets in a market during our sample.

Three variables in the demand model are likely to be endogenous: price, inventory, and the share of the title within its group ($s_{jmt/gmt}$). Since the demand model includes store, title and month fixed effects we are concerned about endogeneity only through unobservables that change over time in one store or title differently from others and that affect changes in prices, inventory and the $s_{jmt/gmt}$ term. We instrument for inventory using the average inventory of the same title across stores of the same tier.[27] Two assumptions are needed to make this a valid instrument. First, that the costs of taking a particular title are correlated across stores of a particular size, implying that similar-sized stores make similar inventory choices. Second, we assume that demand shocks, except those that are captured by the fixed effects δ_j and $\gamma_j z_m$, are not correlated across markets. Examples of shocks captured by the fixed effects are title-specific unobserved advertising and particularly high advertising for a title in college towns. Demand shocks that remain in ξ_{jmt} and that are not correlated across markets could include local advertising and word-of-mouth buzz.

We instrument for $s_{jmt/gmt}$ using two variables: the log of the average number of movies of the same type (same box-genre-store group) in the month, where the average is across other stores in the same size tier that offer the relevant title, and the average of $\ln(s_{jmt/gmt})$ for the same title-month pair across stores of the same tier. The former instrument is correlated with the number of competitors to this title in this store. We take an average over other same-tier stores to account for any demand shocks that might affect both the store's portfolio choice and demand for title j. The second instrument is clearly correlated with $\ln(s_{jmt/gmt})$: like the inventory instrument, it is valid under the assumption that demand shocks, which might affect the share variable, are not correlated across markets.

We tried several instruments for price, including measures of variable costs and average prices of other similar titles. None of the instruments affected the estimate of the price coefficient. We believe the reason for this is that after including store, title and month fixed effects, the only

extended it to include all the stores that actually exist in the market assuming that those we did not observe were identical to those in the data. We would then have simulated the change in demand for store m's titles when all stores simultaneously changed their contracts and portfolios, probably assuming a symmetric equilibrium. Given the lack of data we expect this to add little to our results.

[27] Tiers are defined by Rentrak for the purpose of defining stores' max and min quantity requirements. We assume that they are exogenous to the demand equation modeled here. In all cases we take advantage of the full variation in the data by taking averages over stores in all regions, even when the demand model is run separately for different regions.

unobservable we need to instrument for is at the store-title level, and our potential instruments are not correlated with this price variation. Variation in price at this level (ie., across months) does, however, exist: for example, after a title has been stocked at a store for several weeks, the store may remove the "new release" sticker from the tape and either drop the price or increase the rental period (implying lower collected late fees and a lower observed price). We believe this source of price variation is primarily determined exogenously because of the use of rule-of-thumb policies by video retailers in how they instruct employees to move tapes and update stickers on rental inventory. We therefore conduct our analysis without instrumenting for price. We report in Section 6.2 the OLS results and those that instrument for inventory and the $\ln(s_{jmt/gmt})$ variable.

It is worth noting here that there were other potential demand methodologies. We would ideally have interacted title and store fixed effects in the nested logit; unfortunately the number of parameters to be estimated would then have been infeasibly large. Alternatively we could have estimated a random coefficients model. However, this would have implied replacing the (title or store) fixed effects in the model with (title or store) characteristics. The characteristics available to us are not sufficiently informative for this to be a useful approach.[28]

6.2 Demand Results

We ran the demand model separately for 15 different geographic regions of the country since the dataset was too large for us to run the model using all the data together. We report results for the first geographic region in Table 5.[29] The specification also includes title and store fixed effects and interactions between title fixed effects and store characteristics (percent of the market who are white, percent single and percent with children) and between month fixed effects and title characteristics (box office category, genre and rating and interactions between these). Column 1 of the Table reports results for the OLS regression. Column 2 adds instruments for within-group share and Column 3 also instruments for inventory.

The R^2 is approximately 0.80 in all three models. The fact that the model fits the data well is particularly useful since our supply side estimation will stay within-sample in terms of titles and stores, allowing stores to deviate only in terms of contract choices. We will therefore use all the estimated fixed effects in our inequalitites and counterfactuals.

The price coefficient in the OLS regression is negative and significant, although small. We believe that the size of the coefficient reflects the fact that the demand model captures very short-run demand: rentals at a particular store in a particular month (selected from a set of recently-released titles). We suspect that consumers are not terribly price sensitive in the very short run. For example, even in the event that a consumer's most-preferred title is out of stock at a particular time, she may be likely to rent a substitute title from the same store at a similar price. This leads to low price elasticities in our estimation. In the long-run, of course, the consumer may switch to a

[28] In the next section we regress the sum of the estimated title and store fixed effects on characteristics: the highest R^2 was 0.43.

[29] This region contains zip codes from 20000 to 24999. It includes areas in the mid-Atlantic such as Washington DC and parts of Virginia.

different store; thus, the potential for a store's customer base to shrink in the long-run keeps prices low. If one were to measure this long-run price elasticity, we suspect it would be higher.

The estimated decay rates are intuitive: month 2 demand is higher on average than that in month 1 because observed revenues are left-truncated in month 1 for titles released mid-month. Demand falls in months 3 and 4 and rises again in month 5 because this last observation also includes all subsequent revenues from the title. The inventory coefficient is positive implying that first-month demand increases with the number of tapes on the shelf. Not surprisingly, this generates a reduction in demand in later months (because residual demand is lower). The coefficient on within-group share, σ, is approximately 0.63.

Instrumenting for within-group share reduces the σ coefficient. This is consistent with the existence of demand shocks that affect both within-group share and total demand. Adding instruments for inventory reduces the coefficient on inventory and leaves the interactions between inventory and the decay rate essentially unchanged. There are two potential endogeneity stories here. First, if demand is expected to be high for a particular title then stores will choose high inventory levels, implying a positive bias on all inventory coefficients. Second, heavy advertising of a title in month 1 might lead stores to expect consumers to be impatient, demanding access to the title in month 1 rather than in later months. In this case, the unobservable would lead to high inventory levels and to high demand in the first month; instrumenting should reduce the inventory-month interactions for month 1 only. The results are consistent with the second intuition.

We repeated the demand analysis for each of the 14 other regions. The price coefficient was negative as expected for all but one of these. The problematic region contains 446 stores in the South West, including for example parts of AZ and NV. We exclude this set of stores from the remainder of our analysis.[30] There is some variation in results across the remaining 14 regions. The (unweighted) average price coefficient is -0.026. Its standard deviation across regions is 0.014; the minimum value is -0.050 and the maximum is -0.002. The average inventory coefficient is 0.014, standard deviation 0.004, minimum 0.007 and maximum 0.021. The R^2s in the regressions range from 0.755 to 0.791.

Table 6 sets out the price and inventory elasticities of demand that are implied by our estimates. We calculate the relevant elasticities for each store-title-month triple and then take averages over the observations in each region and each month since release. We then take an unweighted average over regions to generate our summary data. The average elasticity with respect to price over all months since release is low at -0.134. The average elasticity with respect to inventory is 0.155. As for the main demand results, there is some variation across regions. The standard deviation in price elasticities across regions is 0.081; the minimum value is -0.260 and the maximum is -0.011. The standard deviation in inventory elasticities is 0.034 with a minimum of 0.126 and a maximum of 0.242.

Table 6 also documents the variation in elasticities both across months since release of the title and across Box Office categories. There is very little variation over time for the price elasticities.[31]

[30] We also exclude this region from the summary statistics and reduced form analyses in Sections 3 and 4.

[31] Since there is only a single price coefficient, variation in elasticities across months and box-office groups is

The value in month 3 is -0.139, the minimum is -0.143 and the maximum is -0.127. It is worth noting that these elasticities are generated only by price variation that is not absorbed by the store, title and month fixed effects included in our model. It is perhaps not surprising that consumers exhibit only small responses to this within-store price variation, as this in-sample price variation tends to be very small. A more typical methodology would incorporate the effect of cross-store price variation in the elasticity estimates; this is difficult to implement in our application since we rarely observe more than one store per market.[32]

The variation for inventory elasticities is displayed in Figure 1. They are all positive, as expected, implying that a title with more tapes on the shelf has higher demand than other titles. The inventory elasticity for Box Office A titles is higher than those for other Box Office categories, and for each group the elasticity decreases over time since release. These results are consistent with the idea that consumers who rent new releases (particularly Box Office hits) are influenced by advertising such as window displays and the number of tapes available. Consumers who rent movies after the first month or so since release do not expect so much buzz and are less influenced by these kinds of advertising.

Table 7 sets out the results of a regression of the store-title quality levels estimated in the nested logit on store and title characteristics. Our dependent variable is the estimated value of $\left[\hat{\delta}_j + \hat{\gamma}_j z_m + \hat{\eta}_m + \hat{\theta}_t + \hat{\beta}_t x_j + \hat{\lambda}_t c_{jm}\right]$. Results are again reported for the first region; those for other regions are qualitatively similar. The independent variables are title characteristics (quarter of release to video, box office category, genre, rating and interactions of these variables), store characteristics (demographics of the market, the number of households in the market and an indicator for markets where Blockbuster Video is active), interactions between title and store characteristics and the same month dummies and interaction terms that were included in the nested logit. The goal is two-fold: first to check that title and store characteristics have the expected signs, and second to demonstrate the inability of these characteristics to explain the majority of variation in the data.

The results are intuitive. Box office category A titles have higher estimated quality than those in categories B and (particularly) C. Action/adventure movies (the omitted genre category) and comedies have higher demand than other genres; children's movies, romances and science fiction movies have particularly low rental demand. PG13 movies have higher demand than those with other ratings. Markets with a high percent female consumers have high demand for video rentals; those with a high proportion of family heads who are single mothers have particularly high demand and those with a high proportion of family heads who are single without children have particularly low demand. The Blockbuster dummy is positive and significant, probably indicating that Blockbuster chooses to enter high-demand markets.[33] The R^2 on these regressions is only 0.4: even with

generated from variation in the level of demand and price across months and box office categories.

[32]We could regress the estimated store fixed effects on average price and other store characteristics, instrumenting for price with store-level costs. However, most of the variation identifying the coefficients would be across markets rather than across stores.

[33]The coefficient on the number of households is negative and significant, implying that large markets where Blockbuster is not located have low demand. The coefficient on median income is negative and significant: wealthier

a very flexible functional form, our title and store characteristics are able to explain less than half of the variation in the data. This is the reason for using a nested logit framework, in which we can feasibly include both title and store characteristics, rather than adopting a random coefficients model.

7 The Supply Side: Moment Inequalities

Having estimated a detailed demand model, the final piece of information needed to analyze stores' choices of contract types is the cost of holding additional tapes. As noted in Section 5, this includes rent, insurance, restocking costs, the potential value of selling used tapes and of drawing new customers into the store, and also the effect of an increase in inventory of one title on later rentals of other titles that is not incorporated in the demand model. Together these may imply either a positive or a negative estimated average cost of taking additional tapes.

We use the method of moments inequalities estimator developed in Pakes, Porter, Ho and Ishii (2007) to estimate inventory holding costs. That paper describes how to use inequality constraints resulting from a Bayes-Nash equilibrium assumption in both single-agent and multiple-agent games to generate conditions that can be used for estimation and inference. The intuition in our case is very simple: we assume that each store's profit from its observed portfolio of titles and choice of inventory per title must be greater than its profit from any of its alternative choices. We use this assumption to write down a series of inequality constraints. The demand specification models the change in the number of rentals caused by the change in inventory holdings, prices and the consumer's choice set that result from a portfolio deviation. This, together with the price change and change in the number of tapes purchased and the purchase price, will determine the main input to the inequalities analysis: the profit change up to the inventory holding cost.

We choose to use the inequalities methodology in preference to a more standard method such as a logit or probit for two reasons. First, in reality the store is making a dynamic choice. Because each title is active for five months, the store's contract choice for a title released in month 1 should affect its choices of contracts for titles released in months 2-5. For this reason logit and probit methods, which ignore dynamic effects, would generate biased estimates. As discussed below the inequalities methodology enables us to difference out dynamic effects without generating a bias in estimation. The reason is that the alternative portfolio choices defined in the inequalities estimator are not required to be optimal for the store. Second, since the store makes choices regarding FLF contracts at the distributor-year level but those regarding RS and LP at the title level, it would not be obvious at which level to define the logit or probit estimator. No such problem exists for the inequalities methodology.

We derive inequalities from every store's choice of contract for every title.[34] In cases where we observe that the store takes a particular title, we define its alternative choice as dropping that title. This will generate an upper bound on the cost of holding inventory. Where a store does not take a

markets have fewer movie rentals, perhaps because wealthy families choose more expensive leisure activities.

[34]This excludes sell-through priced titles, for which there is no contract choice.

particular title we define its alternative as adding that title on a RS contract. This provides a lower bound on the inventory holding cost. We assume that the store has perfect foresight regarding the titles to be released to video in the following five-month period,[35] but that it may imperfectly predict demand for those titles or the cost of holding inventory. Our assumption, noted above, that each store operates in a separate market implies that no further informational assumptions are needed. We can now write down inequalities that are sufficient to place bounds on the inventory holding cost.

7.1 The Store Profit Equation

Our first step is to predict the total return to the store from its contracts with all distributors over the four year period covered by the data. First, we use the estimated coefficients from the demand model to predict the market share of each title for each store in the market:

$$s_{jmt}(\hat{\delta},\hat{\gamma},\hat{\eta},\hat{\theta},\hat{\lambda},\hat{\alpha},\hat{\sigma},\hat{\xi}) = \frac{e^{(\hat{\delta}_j+\hat{\gamma}_j z_m+\hat{\eta}_m+\hat{\theta}_t+\hat{\beta}_t x_j+\hat{\lambda}_t c_{jm}-\hat{\alpha}p_{jmt}+\hat{\xi}_{jmt})/(1-\hat{\sigma})}}{D_{gmt}^{\hat{\sigma}}\left[\sum_{gmt} D_{gmt}^{(1-\hat{\sigma})}\right]} \tag{4}$$

where:

$$D_{gmt} = \sum_{k\in J_{gmt}} e^{(\hat{\delta}_k+\hat{\gamma}_k z_m+\hat{\eta}_m+\hat{\theta}_t+\hat{\beta}_t x_k+\hat{\lambda}_t c_{jm}-\hat{\alpha}p_{kmt}+\hat{\xi}_j)/(1-\hat{\sigma})} \tag{5}$$

for J_{gmt} the set of all products in group g that are held by this particular store m in month t (other stores are excluded under the assumption that each store essentially operates in an independent market).

Next we consider the return to the store for each title: this is the revenue earned throughout the months after its release less the total payment to the distributor. We denote the return from title j under the three contract types as follows:

1. Under linear pricing the return for title j is $r_{jm}(.) = \sum_{t=t_j}^{t_j+4} q_{tjm}p_{tjm} - F_j c_{jm}$. Here c_{jm} is the capacity of the title (the number of tapes purchased), q_{tjm} is the number of rentals and t indexes time (in months) since the release date t_j.

2. Under revenue sharing the return for title j is $r_{jm}(.) = y_j^{RS}\sum_{t=t_j}^{t_j+4} q_{tjm}p_{tjm} - u_j^{RS}c_{jm}$, where y_j^{RS} is the portion of revenues kept under RS.

3. If the store chooses a full-line forcing contract it has to buy all titles released by the distributor during the following twelve months. It receives better terms than those under revenue sharing: $r_{jm}(.) = y_j^{FLF}\sum_{t=t_j}^{t_j+4} q_{tjm}p_{tjm} - u_j^{FLF}c_{jm}$. Thus $u_j^{FLF} \leq u_j^{RS}$ and $y_j^{FLF} \geq y_j^{RS}$.

We also need to model capacity constraints and quantity restrictions. The number of rentals is constrained by the inventory of the title, c_{jm}, and the maximum feasible number of rentals per

[35] This assumption may well be reasonable: titles' box office release dates are in general more than five months before their release to video so store managers observe which titles will be available several months in advance.

tape, τ_{jm}.[36] Additional constraints, in the form of minimum and maximum quantity restrictions on inventory purchases, are also set by the distributor for RS and FLF contracts (at the store-title level). We denote these constraints as $\underline{c}_{jm}$ and $\bar{c}_{jm}$ respectively. The quantity that would be rented out in the absence of quantity restrictions is:

$$\hat{q}_{jmt} = \min\left(Ms_{jmt}(.), \tau_{jm}c_{jm}\right) \tag{6}$$

The quantity actually rented out is given by:

$$\tilde{q}_{jmt} = \min\left(Ms_{jmt}(.), \tau_{jm}\tilde{c}_{jm}\right) \tag{7}$$

where

$$\tilde{c}_{jm} = \max\left(\underline{c}_{jm}, \min\left(c_{jm}, \bar{c}_{jm}\right)\right) \tag{8}$$

accounts for the effect of the quantity restrictions.

The above implies that the return to the store from a particular title, over the four-year period covered by the data, is given by:

$$\begin{aligned} r_{jm}^{obs}(.) \;=\; & I_{jm}^{LP}\left(\sum_{\tilde{t}=t_j}^{t_j+4}\tilde{q}_{\tilde{t}jm}p_{\tilde{t}jm} - F_j\tilde{c}_{jm}\right) \\ & + I_{jm}^{RS}\left(y_j^{RS}\sum_{\tilde{t}=t_j}^{t_j+4}\tilde{q}_{\tilde{t}jm}p_{\tilde{t}jm} - u_j^{RS}\tilde{c}_{jm}\right) \\ & + I_{jm}^{FLF}\left(y_j^{FLF}\sum_{\tilde{t}=t_j}^{t_j+4}\tilde{q}_{\tilde{t}jm}p_{\tilde{t}jm} - u_j^{FLF}\tilde{c}_{jm}\right) \end{aligned} \tag{9}$$

where time is measured in months. As before we consider the first four months of the lifetime of each title plus a fifth observation for months 5 and above. The indicator functions I_{jm}^{k} equal 1 if contract k is chosen and 0 otherwise.

Given $r_{jm}(.)$, we can write the store's profit from its observed contracts as:

$$\begin{aligned} \pi_m^{obs}(.) \;=\; & \sum_s\sum_{j\in J_s}\left(r_{jm}^{obs}(F,u,y,\bar{c},\underline{c},\hat{\delta},\hat{\gamma},\hat{\eta},\hat{\theta},\hat{\lambda},\hat{\alpha},\hat{\sigma},\hat{\xi},\tilde{c},k) - C(x_m,x_j,\mu)\;\tilde{c}_{jm}\right) \\ & + \eta_m + \rho(\tilde{c}_{ms},k_{ms}) + \varepsilon_{ms} \end{aligned} \tag{10}$$

where J_s is the set of titles released by distributor s during the time period covered by our data and $C(x_m, x_j, \mu)$ is the inventory cost of holding each tape. We estimate this as a reduced form

[36] This is a statement about the technology in use: the maximum number of rentals per title is constrained by the number of times a tape can be rented out before it breaks. In our analysis we define τ to be the 95th percentile of the distribution of the number of rentals per tape that the store provides for titles in the same box office category and contract type in the same month after release. We could add more structure here, accounting for the fact that if a tape breaks in one month it will be unavailable to consumers in later months, but expect this to be a second-order effect in our analysis.

function of store and title characteristics; μ is the parameter vector to be estimated. When the store chooses not to stock a title we assume it makes no contribution to the store's profit. A store fixed effect is denoted η_m, k_{ms} is the contract type (a vector with one element per title for this store-distributor couple) and $\rho(\tilde{c}_{ms}, k_{ms})$ is the effect of the store's choice of contracts at the end of the four-year period on its profits after that period. We choose alternative portfolios whose end-of-period contracts are the same as those observed in the data. The term capturing dynamic effects, $\rho(\tilde{c}_{ms}, k_{ms})$, like η_m, is therefore precisely differenced out of our inequalities.[37] The final term, ε_{ms}, is an unobservable such as store prediction error in the inputs to demand or in the cost of holding inventory.

We account for store capacity constraints by introducing a cut-off condition on inventory. If a particular contract would require the store to take more inventory than a cut-off defined as 110% of the maximum inventory ever taken by the store in any month in our data, then the store adjusts the inventory so that the total number of tapes taken falls below the cutoff. Stores therefore cannot exceed their maximum capacity levels; below those levels all stores' inventory holding costs are the same. As noted above, we do not expect stores in practice to reach their maximum capacity levels.

7.2 The Inequality Estimator

Now consider the store's portfolio choice. As noted above we consider two types of deviations from the store's observed choices. If the store in reality took a title we allow it to drop that title. If in reality it did not take a title we add that title on a RS contract (provided RS was offered for that title).[38] We now demonstrate how the inequalities are generated in these cases.

Take as an example a title t' released by distributor s'. Suppose that the store chose to take the title on a LP contract. We assume that:

$$E\left(\pi_m^{obs}(.) \mid I_m\right) \geq E\left(\pi_m^{altt'}(.) \mid I_m\right) \tag{11}$$

for the observed portfolio choice, where $\pi_m^{altt'}$ is defined analogously to π_m^{obs} but using the alternative portfolio choice where the store drops title t' and holds all other contracts fixed. The expectation is taken conditional on I_m, the store's information set at the time when it makes its choice. We therefore infer from the observed data that:

$$E\left(\pi_m(k_{ms'}^{LP}(t')) \mid I_m\right) \geq E\left(\pi_m(k_{ms'}^{0}(t')) \mid I_m\right) \tag{12}$$

where $k_{ms'}^{LP}(t')$ indicates that the t'th element of $k_{ms'}$ is a LP contract and $k_{ms'}^{0}(t')$ indicates that the title has been dropped. This equation implies the following inequality (where title t' has zero demand by the end of the 4-year period in the data and therefore the $\rho(.)$ term is differenced out

[37] This modeling choice may imply that the alternative portfolio used in our inequalities is not the optimal alternative for the store. However, as discussed below this will not affect the consistency of the estimates.

[38] We consider only titles that are in release in our data for at least six months.

along with the store fixed effect):

$$E\left(\Delta\pi_m^{s',t'}(.) \mid I_m\right) =$$

$$E\{\sum_s[\sum_{j\in J_s}\{\Delta r_{jm}(k_{ms'}^{LP}(t'),k_{ms'}^0(t')) - C(.)\ \Delta\tilde{c}_{jm}(k_{ms'}^{LP}(t'),k_{ms'}^0(t'))\} + \Delta\varepsilon_{ms'}] \mid I_m\} \geq 0 \qquad (13)$$

Here the difference function $\Delta\tilde{c}_{jm}(k_{ms'}^{LP}(t'),k_{ms'}^0(t')) = \tilde{c}_{jm}^{obs} - \tilde{c}_{jm}^{altt'} = \tilde{c}_{jm}(k_{ms'}^{LP}(t')) - \tilde{c}_{jm}(k_{ms'}^0(t'))$, and similarly for $\Delta r(.)$. The calculation incorporates the revenue from all titles released by all distributors in all years in the data, since changing a single contract may affect demand for other-distributor titles, even if these are released in later months.[39] The returns from both the observed and the alternative portfolios are calculated from the model to ensure comparability in the counterfactual. Each title taken by the store implies a similar inequality. We consider a switch from not taking a title to taking it on RS in an analogous way.

Our next step is to take an expectation conditional on the instruments $z_{ms'}$, where s' is the distributor whose titles were switched by the store. We define these instruments such that $z_{ms'} \subset I_m$ and $E(\varepsilon_{ms'} \mid z_{ms'}) = 0$. This together with equation (13) implies that:

$$E(\Delta\pi_m^{s',t'}(.) \mid z_{ms'}) = E\left\{\sum_s\sum_{j\in J_s}\left(\Delta r_{jm}^{s',t'}(.) - C(x_m,x_j,\mu)\ \Delta\tilde{c}_{jm}^{s',t'}(.)\right) \mid z_{ms'}\right\} \qquad (14)$$

$$= E\left\{\Delta r_m^{s',t'}(.) - C(.)\Delta c_m^{s',t'}(.) \mid z_{ms'}\right\} \geq 0 \qquad (15)$$

where $\Delta r_m^{s',t'}$ is the sum of $\Delta r_{jm}^{s',t'}$ over all titles and the change in total store-level inventory holding cost for a particular deviation is written as:

$$C(.)\ \Delta\tilde{c}_m^{s',t'}(.) = \sum_s\sum_{j\in J_s} C(x_m,x_j,\mu)\ \Delta\tilde{c}_{jm}^{s',t'}(.).$$

The variables in $\pi_m(.)$ that will change under the alternative portfolio are the inventory c_{jm}, prices p_{jtm}, technology (rentals per tape) τ_{jm}, and contract terms. Our assumptions regarding prices and quantities are as follows. We note that prices vary only slightly between titles within a store. The average mean within-store price of an A title is \$2.88 for RS contracts and \$2.84 for LP. The equivalent prices for B titles are \$2.79 and \$2.80; those for C titles are \$2.73 and \$2.73 respectively. The variation is even smaller within contract group. We therefore do not directly model a price change after the change in portfolio. Instead we use the average price for each month for the particular store-box office category-contract type-month combination being considered.

Similarly, we do not formally model the firm's choice of c_{jm} and τ_{jm} for every title. We define the quantity $\tilde{q}_{tjm}$ as in equation (7). The first term, $Ms_{tjm}(.)$, represents consumer demand for the

[39] In fact the demand framework only allows a change in contract for title j to affect the within-group share and therefore demand for title k in months where they overlap in consumers' choice sets. It seems reasonable to assume that title k's demand in months before j is released will be unaffected by a change in j's contract type, assuming that consumers do not predict this change. If title k is active after title j has left the dataset, we assume that its demand in these later months is unaffected by j's contract change.

title in month t. We predict this using the estimated demand coefficients, the other titles offered by the store, and the expected price and inventory choices for the relevant titles, defined as averages over other titles in the same store-box office category-contract type-month. The inventory level is also constrained by the maximum and minimum quantity restrictions for that title as defined in equation (8).[40] The last term, $\tau_{jm}\tilde{c}_{jm}$, is the maximum number of rentals the store can offer for this title given the contract type. We interpret this as the store's inventory level for the title under the relevant contract type multiplied by its maximum τ_{jm} (the maximum number of rentals per tape). This maximum τ_{jm} is defined as the 95th percentile of the distribution of τ observed for titles in the same store-box office category-contract type. The inventory level is the same average used as an input into expected demand.[41,42] We define the contract terms of the alternative contracts (the RS contracts in the cases where a new title is added) using the modal values over all stores for that contract type and title.[43]

It is worth noting here the distinction between the different methods of forming expectations used in our analysis. When stores choose their contract types, the contract terms (split, upfront cost, wholesale price and maximum/minimum quantity restrictions) are defined by the distributor and perfectly observed by the store. We predict expected prices using averages across titles within the same store-box office category-contract type-month since there is less variation within-store across titles than there is across stores for a particular title (see Mortimer (2007) for evidence on this).[44] Expected inventory is treated analogously, except that it does not vary by month. Where a title is observed to have zero transactions in a particular month at a store, we follow a methodology consistent with that used in the demand model and exclude these observations.

[40] We define these values as the modal values for that title and contract type across same-tier stores. If the relevant title-contract-tier group is empty we fill in values using neighboring tiers.

[41] When the title is newly taken under a RS contract as a deviation from the observed store choice, we use the minimum number of tapes required by the distributor as the inventory level rather than the average because these titles are likely to have lower quality than average and therefore to be in relatively low demand from consumers. If taking a new title would force the store to hold a total storewide inventory level higher than 110% of that observed for the store in any month in the data, we assume that the store's inventory would equal that maximum cutoff, provided it implies an inventory level for the new title that is above the minimum quantity restriction of the contract; otherwise no inequality is generated.

[42] If the store-box-contract group is ever empty we use the minimum number of tapes required by the distributor as the expected inventory level for RS and FLF titles. For LP titles we use the average number of tapes per LP title taken by the store for titles in neighboring box office categories. Remaining missing values for LP contracts are filled in using $\frac{1}{3}$ of the minimum quantity required under RS for the same title: this is the industry rule of thumb for the number of tapes taken under LP. We predict price and τ_{jm} in the case where the store-box-contract group is empty using other contract types. Prices for titles taken on LP or STP contracts are defined as the maximum of the average prices in the same store-box group under RS, FLF and STP contracts; those for titles taken on RS or FLF contracts (the latter where the store is observed to take the title under a different contract type) are the average of the average prices under the other three contract types. Finally, in all these cases we define maximum τ as the maximum of the maximum τ's under the other three contract types. If any of these values are still missing we use the average price or maximum τ in the store-contract type for titles in neighboring box office groups or the average price in the store-contract type-box office category for neighboring months.

[43] In fact they are constrained by Section 2 of the Clayton Act to be the same for all stores for a particular title. We take a modal value because a small number of stores negotiate special deals such as volume discounts with particular distributors. These are classed as second-degree price discrimination and are therefore not illegal. We assume that stores do not expect to be able to negotiate such deals for alternative contracts.

[44] We could have considered titles in the same store-genre-box office category-contract type-month, consistent with the nest definitions in the demand model, but encountered problems with small sample sizes in some cases.

Finally we convert the expectations in equation (14) into sample averages across stores.[45] We also average over alternative choices t' in a particular distributor-year before interacting with the instruments. Considering first the inequalities for titles that are observed to be taken by the relevant store, this implies the following equation for estimation:

$$\Delta\bar{\pi}_{ys'} = \frac{1}{M}\sum_{m}\left((r_m^{obs} - C(.)\ \tilde{c}_m^{obs}) - \frac{1}{Q_{ys'}}\sum_{t'\in(s',y)}(r_m^{altt'} - C(.)\ \tilde{c}_m^{altt'})\right) \otimes g(z_{ms'}) \geq 0 \qquad (16)$$

where y indexes years, s' indexes distributors, $Q_{ys'}$ is the number of titles released by distributor s' in year y, $g(.)$ is any positive-valued function of the instruments and M is the number of stores in the data. We therefore have one moment per distributor-year-instrument triple. We exclude distributor-years where no titles are released and average within each year (before taking the store average) over distributors that released fewer than five titles during our panel. This generates 56 distributor-year moments per instrument. Our methodology is very similar when the switch is from not taking the contract to taking it on RS, except that we include only titles that were offered on RS contracts, generating 38 additional distributor-year moments per instrument.[46] The identified set of parameter values is the set of parameters that satisfy the implied system of inequalities. If there are no feasible parameters we use a method of moments methodology, minimizing the Euclidean distance by which the inequalities are violated.

In each inequality we hold the portfolio of titles fixed other than the single change being considered. Of course in reality the store may change both its portfolio of titles offered by this distributor and the set of titles taken from other distributors when it makes that single change. However, we do not need to model the store's optimal portfolio choices here in order to consistently estimate the inventory holding cost. The simpler inequalities that hold the rest of the portfolio fixed are also valid (since they are generated from feasible alternatives for the store) and are sufficient for our purposes. This simplification has the added benefit that it enables us to difference out the dynamic effects captured by the $\rho(.)$ term in equation (10). We simply choose not to change the observed contracts for store-title pairs for which we observe less than five full months of data, that is, those whose direct influence on other titles' revenues or optimal contract types extend beyond the end of our panel.[47] We do, however, model portfolio choices in the counterfactual analyses considered below.

The instruments $z_{ms'}$ (defined at the store-distributor level) are required to be uncorrelated

[45] We stack the inequalities for all regions before taking the average over stores. We therefore estimate a single set of costs for all stores, taking into account all of the region-specific demand estimates from Section 6.

[46] We define the unobserved quality of the new titles to be the minimum ξ_{jmt} in the store-box-genre-month group. If there are no other titles in this group we use the average in the store-box-month. If no stores took the title in the relevant zip code category we have no δ_{jmt} estimates for that region; we exclude these zipcodes from the counterfactuals for the relevant title. Any unobserved differences between these zipcode categories and the others that might affect the results will be captured in the store fixed effects in the store profit equation.

[47] The single title whose contract is changed in the alternative scenario may indirectly influence the optimal contract choice of later titles since it overlaps in time with other titles whose direct influence may extend beyond our panel. However, again we can ignore this effect since the alternative considered need not be optimal for the store.

with $\varepsilon_{ms'}$, the unobservable in the profit equation, and correlated with the capacity chosen by the store. The unobservable includes store prediction error regarding demand for particular titles. It could include variation in inventory holding costs and other store costs that are not observed by the store but that will affect its total costs when it alters its contracts with distributor s'. It could also include store prediction error regarding the prices for which the store will be able to sell used tapes or of the proportion of tapes that will break before they can be sold. Our instruments include the number of titles released by the distributor in the relevant year and the percent of these titles that are from Box Office categories A and B respectively. At the store level we use indicators for stores with a high percent single population, a high percent of the population with children, a high median age and a high number of households, where "high" is defined as above the 75^{th} percentile in the data. We also use indicators for the size of the store's chain.

The inequality method will lead to biased estimates if the unobservable ε_{ms} contains any variables that differ across contract types and are observed by the store. For example, there may be differences in the restrictions placed by distributors on stores' sales of used tapes for LP contracts compared to RS and FLF.[48] We do not accurately observe the contract-specific requirements and therefore cannot control for them in our estimation. It is reassuring to note that sales of used tapes make up a fairly low proportion of each store's revenues, especially for Box Office category B and C titles. We might also be concerned about the endogeneity of $\bar{c}_{jm}$ and $\underline{c}_{jm}$, both of which are set by the distributor on a title-by-title basis. Any unobservable that affects the store's choice of contract for title j may also affect the distributor's choice of quantity restrictions. However, many distributors choose these quantity restrictions using a formula based on the title's box office sales and the size of the store.

7.3 Results

Our specification for the cost of holding a tape includes three variables: a constant and indicators for Box Office category B and C titles. The results are given in Table 8. The estimate of μ was a singleton: that is, there was no parameter vector that satisfied all the inequality constraints.[49] We report the conservative 95% confidence intervals derived in Pakes, Porter, Ho and Ishii (2007). We estimate a negative cost of \$11.30 of holding each tape of Box Office category A titles (those with theatrical box office revenues over \$40 million). The "value" of B titles (those with box office revenues of \$15-40 million) is essentially zero. C titles (those with box office revenues under \$15 million) are estimated to have a positive value of \$11.80 per tape.

There are several potential sources of a positive value per tape to the store. First, the store generates approximately \$9 of revenues from selling each used tape. Second, unobserved volume

[48] A title that is obtained on a LP contract will be sold at a price set by the store at the end of its rental life; the store retains 100% of the revenues from these sales. However, if the store obtains a title on a RS or FLF contract it is required to pay a certain proportion of the revenues from used tape sales back to the distributor.

[49] As noted in Pakes, Porter, Ho and Ishii (2007), this does not imply that we should reject the specification. The result could easily be caused by the random disturbances in the inequalities. The probability that all the inequalities will be satisfied can be made arbitrarily small by increasing the number of inequality restrictions.

discounts, particularly under LP contracts, could reduce the cost to the store of taking tapes compared to the cost used in our analysis. Conversations with industry experts indicate that these discounts are probably approximately 20% of the wholesale price for LP titles, or $13 per tape on average. Third, adding tapes to the shelf (particularly those for titles in Box Office category A) may increase retailer profits by attracting new consumers to the store. For all of these reasons the small estimated value to the store of A and C titles seems quite reasonable.

We estimate that B titles have a zero value per tape. The reason why their value is lower than that of C titles may be that C titles are often niche or "arthouse" titles that appeal to a particular segment of the population and may be more attractive to specific consumers - and therefore bring in more revenues to the store from the sale of used tapes and from drawing in customers from the store's competitors - than B titles which are often relatively unsuccessful mainstream movies.

8 Counterfactual Analyses

Our goal in the final stage of the analysis is to conduct a counterfactual that allows us to estimate the welfare effect of FLF contracts. Our informational assumption for estimating the moment inequalities is that each store perfectly predicts the titles to be released in future months but may imperfectly predict consumer demand for these titles and the cost of holding inventory. Our counterfactual experiment is consistent with this assumption.

We use our estimates to predict the effect of postponing the implementation of FLF contracts for a subset of the distributors in our data. Two of the seven FLF distributors in our data implemented FLF early, in month 14 of our panel. The others followed some time later: in or after month 30. We leave the two early implementers unchanged and remove the three earliest months of FLF for the distributor that began FLF in month 30. For each store we consider all possible combinations of which titles to take for the FLF distributor over the three month period in which we delay the introduction of the FLF contract. Our assumption that each store operates in a separate market implies a unique equilibrium. We obtain a set of choices over titles' contract types that are optimal for the store assuming perfect foresight about the titles to be released in that three-month period, given the observed choices in the data both before and after the three-month counterfactual period.

We hold fixed stores' choices of contract types both before and after our counterfactual. We assume that the removal of FLF was a surprise, so choices made before the removal may not be optimal after the change. Choices made after the counterfactual may in fact be affected by it: for example FLF contracts observed in the data in the first month after FLF is re-instated may have started in earlier months, implying that they should be extended in our counterfactual. In addition, the observed LP and RS contracts in the first few months may not be optimal given the new contract choices during the counterfactual period. We use two alternative methods to address this issue. First we assume that stores make their choices in the belief that FLF will be available from the beginning of the counterfactual onwards. The FLF distributors then unexpectedly delay the beginning of the program by three months but hold stores to their commitments to purchases

after FLF is reinstated. Other distributors also hold stores to these commitments but allow them to re-optimize during the counterfactual period. This latter assumption is probably unrealistic. Our second counterfactual instead assumes that the non-FLF distributors hold stores to their prior commitments throughout the counterfactual period and beyond it. This rules out the leverage effect but may not have much impact on the overall results (particularly given our reduced form analyses which indicate that the leverage effect may be small).

In both counterfactuals we compare the model's predictions for store choices, store and distributor profits and consumer surplus given the delay in FLF introduction to the model's predictions without the delay. The "no delay" scenario generates predictions for store contract choices, rather than using the observed contract choices, to ensure comparability across scenarios. We also present a comparison of stores' predicted contract choices to their observed choices under the "no delay" scenario; we view this as a test of the goodness-of-fit of our model. In all cases we allow the store to optimize over the set of titles taken from the relevant distributors and the contract type of these titles. We simplify by assuming, as we did in the moment inequalities analysis, that inventory and rental price are determined by contract type (the average in the data for the store-box-contract type-month). We define the unobserved quality of newly-taken titles to be the 25th percentile in the store-box-genre group.

Our results are displayed in Table 9. Five titles were released by the FLF distributor during the three-month period of our counterfactual. 56 were released by other distributors during the same time period. We consider only stores that we observe to take all five of the distributor's titles under FLF: that is, stores that are directly affected by the counterfactual. There are 144 of these stores: it is feasible to include 116 of them (representing 79% of the 144 stores' revenues) in the counterfactual.[50] Of these 116 stores, 83 are predicted by our model to take FLF. This is another measure of the fit of our model: 72% of stores are predicted to make the choice observed in the data. We exclude the remaining 33 stores from the counterfactual.

Column 1 of Table 9 reports the model's predicted dollar values for mean per-store profits, the total profits of the FLF distributor and of other distributors, and means of per-store consumer surplus and per-store total welfare under the observed store contract choices. Column 2 reports these values using our model's predictions for store contract choices under the "no delay" scenario (our new baseline) and the percent difference from the values in Column 1. The differences are very small, ranging from -0.64% for FLF distributor profits to 1.33% for mean per-store profits. The model fits the data well for the stores included in the counterfactual.

Column 3 of the Table reports the change in outcomes under the counterfactual that excludes the leverage effect. The bottom panel shows that on average 1.8 of the FLF distributor's titles are dropped by the store when FLF is delayed; 0.3 are switched from FLF to LP and 2.9 switched to RS. This is consistent with a fairly significant market coverage effect: approximately 1.8 out of 5 titles were taken under FLF that would not otherwise have been taken by the store. This change

[50]The remaining 28 stores take a very small number of titles from competing distributors, which increases the number of combinations of possible replacement titles that the simulation needs to consider. This increase requires a few months of additional computational time on available computer equipment.

in uptake results in a mean per-store profit decrease for the relevant 3 months of 4.67%. The profit reduction makes sense intuitively: these are stores that chose to take the FLF contract so we should expect their profits to fall when we remove that option. The FLF distributor's profits for the same 3 months, from the stores included in the counterfactual, fall by 30% in total (assuming a cost to the distributor of $2 per tape). Other distributors' profits increase slightly (by just 0.6%) because some consumers switch from the dropped FLF titles, and from those for which inventory has decreased with the change in contract type, to the titles released by other distributors. Consumer surplus falls slightly because of the reduction in the size of the choice set. The total effect of delaying the introduction of FLF is a mean welfare loss of 1.87% per store.

Finally Column 4 of the Table reports the change in outcomes compared to the baseline (Column 2) when we also allow for a leverage effect. It is not feasible to allow the store to consider changing its contracts with all titles released by all distributors during our three month time period. Instead we consider only titles not taken by the relevant store that are in the same box-genre-rating groups as the five titles released by the FLF distributor, since these are the most likely to be added when the store drops FLF titles. The bottom panel of Table 9 shows that the leverage effect is very small: only 0.012 titles are added on average on LP contracts and 0.024 titles added on RS. Not surprisingly, the welfare changes from this counterfactual are very similar to those without the leverage effect. The only difference that is apparent to 2 decimal places is a slightly larger increase in the profits of other distributors. This is due to the small increase in the number of other-distributor titles taken by the store. The overall welfare effect of delaying FLF is a mean welfare loss of 1.85% per store, a little lower than the 1.87% loss when we exclude the leverage effect. The market coverage and efficiency effects clearly dominate the leverage effect in our application.

9 Conclusion

The results from the full model are consistent with the reduced form analysis. Store profits fall slightly from the delay in FLF and the FLF distributor loses a substantial amount. This loss far outweighs the small gain to other distributors. The leverage effect is much smaller than the combined market coverage and efficiency effects. Our final results therefore indicate a positive overall welfare effect of FLF contracts (a loss from delaying their introduction). To some extent this is not surprising. The reasonably large number of active distributors (59 in our data), each producing only on average 8 titles per year, implies that each distributor's power to persuade retailers to exclude other distributors is probably quite limited. The finding of a small leverage effect is therefore not unexpected. The market coverage and efficiency effects are both welfare-improving, implying a positive aggregate effect of FLF contracts on welfare. Further study of the welfare effects of bundling contracts, particularly in settings where the upstream firm may have more market power, would be useful.

References

Adams and Yellen, (1976), "Commodity Bundling and Burden of Monopoly," *The Quarterly Journal of Economics*, V. 90, No. 3, pp. 475-98.

Bork, (1978), *The Antitrust Paradox: A Policy at War with Itself,* New York, Basic Books, Inc.

Burstein, (1960), "A Theory of Full-Line Forcing," *Northwestern University Law Review*, V. 55, pp. 62-95.

Byzalov, D., (2008), "Unbundling Cable Television: An Empirical Investigation," *working paper.*

Carbajo, de Meza and Seidmann, (1990), "A Strategic Motivation for Commodity Bundling," *Journal of Industrial Economics*, V. 38, No. 3, pp. 283-98.

Carlton and Waldman, (2002), "The Strategic Use of Tying to Preserve and Create Market Power in Evolving Industries," *RAND Journal of Economics*, V. 33, No. 2, pp. 194-220.

Cardell, N. Scott, (1997), "Variance Components Structures for the Extreme Value and Logistic Distributions with Applications to Models of Heterogeneity," *Journal of Econometric Theory*, V. 13, pp. 185-213.

Choi and Stefanadis, (2001), "Tying, Investment, and the Dynamic Leverage Theory," *RAND Journal of Economics*, V. 32, No. 1, pp. 52-71.

Chu, C., Leslie, P. and Sorenson, A., (2007), "Nearly Optimal Pricing for Multiproduct Firms," *working paper.*

Conlon, C. and Mortimer, J., (2008), "Demand Estimation Under Incomplete Product Availability," *working paper.*

Crawford, G., (2008), "The Discriminatory Incentives to Bundle in the Cable Television Industry," *Quantitative Marketing and Economics*, V. 6, No. 1, pp. 41-78.

Crawford, G., and Yurukoglu, A., (2008), "The Welfare Effects of Bundling in Multi-Channel Television Markets," *working paper.*

Marx, L. and Shaffer, G., (2004), "Slotting Allowances and Scarce Shelf Space," *working paper.*

McAfee, McMillan and Whinston, (1989), "Multiproduct Monopoly, Commodity Bundling and Correlation of Values," *The Quarterly Journal of Economics*, V. 104, No. 2, pp. 371-83.

Mortimer, J., (2007), "Price Discrimination, Copyright Law and Technological Innovation: Evidence from the Introduction of DVDs," *The Quarterly Journal of Economics*, V. 122, pp. 1307-50.

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

Nalebuff, (2004), "Bundling as an Entry Deterrent," *Quarterly Journal of Economics*, forthcoming.

Shy, Oz, (1995), "Industrial Organization," The MIT Press, Cambridge, MA.

Pakes, Porter, Ho and Ishii, (2007), "Moment Inequalities and Their Application," *working paper.*

Posner, (1975), "The Social Costs of Monopoly and Regulation," *Journal of Political Economy*, V. 83, N. 4, pp. 807-27.

Salinger, (1995), "A Graphical Analysis of Bundling," *Journal of Business*, V. 68, N. 1, pp. 85-98.

Schmalensee, (1982a), "Antitrust and the New Industrial Economics," *American Economic Review, Papers and Proceedings*, V. 72, N. 2, pp. 24-28.

Schmalensee, (1982b), "Commodity Bundling by Single-Product Monopolies," *Journal of Law and Economics*, V. 25, N. 1, pp. 67-71.

Stigler, G., (1962), "United States vs Loew's Inc: A Note on Block-Booking," *The Supreme Court Review*, V. 1963, pp. 152-157.

Whinston, (1990), "Tying, Foreclosure and Exclusion," *The American Economics Review*, V. 80, N. 4, pp. 837-59.

Figure 1: Inventory Elasticities Implied by Demand Estimates

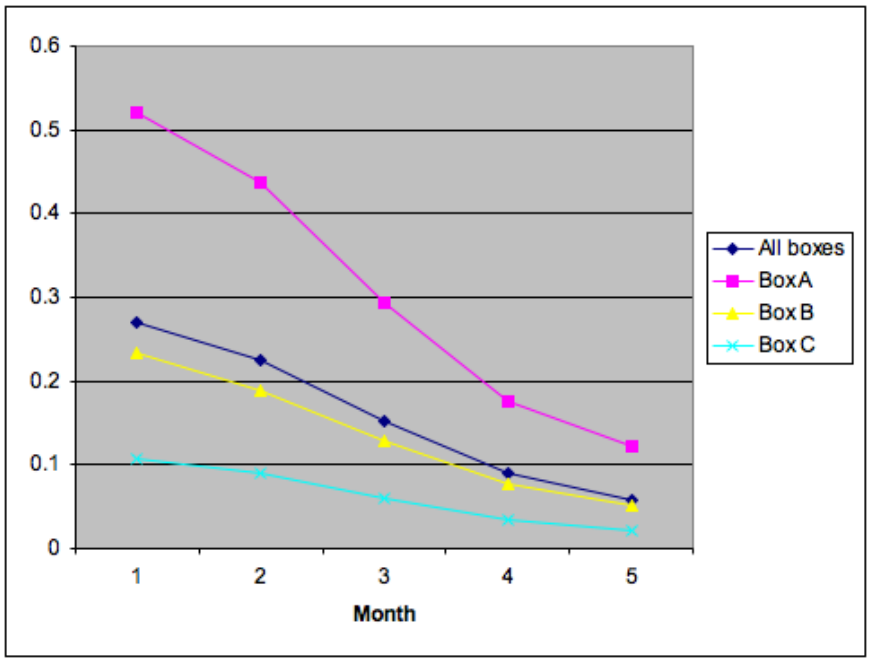

Notes: Elasticities implied by the demand estimates. Corresponding data are given in Table 6. Elasticities with respect to inventory are calculated for every store-title-month observation and then averages are taken within each zipcode region - month and then across regions.

Table 1: Summary Statistics

Contract	Linear Pricing	Revenue Sharing	Full-Line Forcing	Sell-Through Pricing
Avg. Upfront Cost	66.82	8.47	3.62	15.17
	(5.60)	(1.08)	(1.20)	(1.63)
Avg. Retailer's Share of Revenue	100%	45.99%	59.02%	100%
	(-)	(2.99%)	(1.99%)	(-)
Avg. Minimum quantity	-	10.32	10.87	-
	(-)	(11.55)	(10.53)	(-)
Avg. Maximum quantity	-	23.49	22.61	-
	(-)	(22.48)	(21.38)	(-)
Avg No. of Rentals				
Month 1:	52.11	62.58	52.16	91.28
	(81.85)	(86.88)	(88.72)	(120.49)
Month 2:	67.36	66.21	61.87	85.32
	(92.91)	(91.18)	(82.27)	(101.92)
Month 3:	40.22	33.24	33.12	41.31
	(52.88)	(44.99)	(44.23)	(47.93)
Month 4:	25.90	21.22	20.93	23.00
	(32.42)	(27.54)	(26.94)	(25.80)
Month 5+:	69.89	56.79	51.45	74.56
	(101.98)	(83.37)	(74.42)	(134.51)
Avg Rental Price				
Month 1:	2.65	2.68	2.68	2.69
	(0.57)	(0.49)	(0.60)	(0.53)
Month 2:	2.83	2.78	2.88	2.87
	(0.56)	(0.50)	(0.56)	(0.59)
Month 3:	2.83	2.78	2.87	2.94
	(0.61)	(0.55)	(0.64)	(0.71)
Month 4:	2.83	2.78	2.86	2.96
	(0.67)	(0.60)	(0.68)	(0.83)
Month 5+:	2.79	2.68	2.86	2.93
	(0.72)	(0.67)	(0.76)	(0.89)
Avg Rentals per Tape				
Month 1:	5.58	4.19	4.14	5.26
	(4.44)	(2.72)	(3.14)	(4.85)
Month 2:	7.59	4.71	5.54	5.19
	(4.93)	(3.11)	(3.84)	(3.75)
Month 3:	5.06	2.51	3.46	2.71
	(3.83)	(1.86)	(2.91)	(2.17)
Month 4:	3.59	1.67	2.48	1.65
	(2.98)	(1.32)	(2.43)	(1.69)
Month 5+:	13.59	5.00	7.42	6.72
	(14.47)	(4.93)	(9.56)	(8.99)
Avg Inventory	9.09	14.20	12.74	18.18
	(14.25)	(16.86)	(17.27)	(21.60)

Notes: Averages are across store-title pairs. Standard deviations in parentheses.

Table 2: Summary Statistics (Cont.)

Contract	Linear Pricing	Revenue Sharing	Full-Line Forcing	Sell-Through Pricing
Total No. of Titles Released by Distributors				
Year 1:	219	115	0	27
A Titles:	30	12	0	15
B Titles:	36	17	0	6
C Titles:	153	86	0	6
Year 2:	204	125	10	24
A Titles:	32	23	1	14
B Titles:	42	29	2	6
C Titles:	130	73	7	4
Year 3:	231	132	18	21
A Titles:	43	29	4	15
B Titles:	44	29	3	1
C Titles:	144	74	11	5
Year 4:	209	113	38	26
A Titles:	36	20	9	16
B Titles:	50	19	5	3
C Titles:	123	74	24	7

Notes: Total number of titles released by distributors and offered under each contract type. Titles may be counted in more than one column. All Revenue-Sharing and Full-Line Forcing titles are also offered under Linear-Pricing contracts. No Sell-Through Pricing titles are offered under alternate contracts.

Table 3: Summary Statistics (Cont.)

Contract	Linear Pricing	Revenue Sharing	Full-Line Forcing	Sell-Through Pricing
Number of Stores				
	6,358	6,150	5,111	6,171
Avg No. of Titles Taken by Stores				
Year 1:	32.54	7.92	-	6.72
	(23.80)	(11.21)	-	(4.56)
A Titles:	19.57	4.94	-	11.83
	(7.52)	(3.56)	-	(4.12)
B Titles:	23.61	5.29	-	4.27
	(9.33)	(5.39)	-	(1.64)
C Titles:	54.51	13.56	-	3.99
	(28.74)	(16.97)	-	(1.69)
Year 2:	23.15	6.84	1.39	5.05
	(18.97)	(7.74)	(1.53)	(4.11)
A Titles:	14.29	7.04	0.61	8.87
	(8.69)	(5.99)	(0.49)	(4.34)
B Titles:	19.94	6.94	1.22	3.85
	(11.68)	(7.68)	(0.85)	(2.37)
C Titles:	35.63	6.54	2.36	2.32
	(25.23)	(9.25)	(2.14)	(1.55)
Year 3:	32.69	6.30	2.09	4.90
	(23.50)	(8.37)	(2.07)	(5.57)
A Titles:	23.35	7.32	1.57	11.23
	(12.18)	(7.31)	(0.98)	(5.21)
B Titles:	26.26	5.35	0.92	0.78
	(13.20)	(7.51)	(0.66)	(0.41)
C Titles:	48.79	6.20	3.80	2.53
	(30.95)	(9.95)	(2.64)	(1.43)
Year 4:	28.26	4.67	3.36	5.79
	(15.20)	(6.14)	(3.44)	(5.16)
A Titles:	19.64	5.66	3.33	11.74
	(8.65)	(5.30)	(2.07)	(4.78)
B Titles:	32.59	3.69	1.05	2.41
	(12.55)	(5.37)	(0.98)	(1.06)
C Titles:	32.60	4.64	5.72	3.18
	(18.72)	(7.37)	(4.40)	(1.48)

Notes: Average number of titles of each contract type taken by all active stores in each year. Standard deviations in parentheses.

Table 4: Summary Statistics (Cont.)

Contract		Linear Pricing	Revenue Sharing	Full-Line Forcing	Sell-Through Pricing
Ave store tier					
Quintile 1		2.68	5.04	3.77	5.47
Quintile 2		2.95	5.05	2.78	4.65
Quintile 3		3.51	3.63	4.26	3.98
Quintile 4		4.79	2.97	4.25	3.03
Quintile 5		5.46	2.70	4.31	2.27
Quintile	% of quintile				
1	Tier 1-3	1.59	0.43	1.13	0.54
	Tier 7-10	0.53	1.71	1.40	2.99
2	Tier 1-3	1.41	0.41	1.48	0.56
	Tier 7-10	0.31	1.41	0.30	1.20
3	Tier 1-3	1.15	1.19	0.71	0.84
	Tier 7-10	0.65	1.00	0.80	0.54
4	Tier 1-3	0.47	1.40	0.84	1.34
	Tier 7-10	1.11	0.34	1.16	0.15
5	Tier 1-3	0.37	1.57	0.85	1.72
	Tier 7-10	2.39	0.53	1.30	0.12

Notes: Panel 1 breaks the percent of each store's titles adopted under a particular type of contract into quintiles and reports the average store tier in each quintile. Tiers are ranked from 1 to 10 where 10 is largest. Panel 2 reports the percent of stores in each quintile that are in store tiers 1-3 and 7-10 respectively. These percentages are normalized by the percent of all stores that are in the relevant set of tiers. Numbers over 1 indicate that the store type is over-represented in the relevant quintile.

Table 5: Demand Results

	OLS Coefft (S.E.)	IV 1 Coefft (S.E.)	IV 2 Coefft (S.E.)
Price	-0.027 (0.002)	-0.026 (0.003)	-0.024 (0.003)
Month 2	0.132 (0.023)	0.155 (0.024)	0.139 (0.025)
Month 3	-0.136 (0.022)	-0.191 (0.025)	-0.206 (0.025)
Month 4	-0.399 (0.023)	-0.505 (0.025)	-0.512 (0.026)
Month 5+	0.190 (0.024)	0.276 (0.027)	0.283 (0.027)
Inventory	0.019 (0.0003)	0.021 (0.0004)	0.016 (0.0005)
Inv*Month 2	-0.003 (0.0004)	-0.004 (0.0005)	-0.003 (0.0005)
Inv*Month 3	-0.008 (0.0004)	-0.009 (0.0005)	-0.008 (0.0005)
Inv*Month 4	-0.012 (0.0004)	-0.013 (0.0005)	-0.013 (0.0005)
Inv*Month 5	-0.011 (0.0004)	-0.013 (0.0005)	-0.014 (0.0005)
σ	0.632 (0.0018)	0.498 (0.0030)	0.501 (0.0031)
N	407,006	407,006	407,006
R^2	0.82	0.76	0.76

Notes: Results of nested logit demand analysis. IV1 instruments for the within-group share only. IV2 instruments for both within-group share and inventory. All specifications include title and store fixed effects, interactions between title fixed effects and store characteristics (the percent in the market with kids, the percent single and the percent white) and interactions between month fixed effects and title characteristics (the box office category, genre, rating and interactions of these variables).

Table 6: Demand Results: Elasticity Estimates

	Month 1	Month 2	Month 3	Month 4	Month 5
All Box Office categories:					
Price elasticity	-0.127	-0.132	-0.139	-0.143	-0.129
Inventory elasticity	0.269	0.224	0.151	0.090	0.058
Box Office Category A:					
Price elasticity	-0.125	-0.132	-0.142	-0.148	-0.138
Inventory elasticity	0.521	0.437	0.294	0.176	0.123
Box Office Category B:					
Price elasticity	-0.125	-0.127	-0.136	-0.140	-0.128
Inventory elasticity	0.234	0.188	0.128	0.078	0.052
Box Office Category C:					
Price elasticity	-0.129	-0.135	-0.139	-0.141	-0.123
Inventory elasticity	0.106	0.089	0.059	0.035	0.022

Notes: Elasticity estimates implied by the demand estimates. Demand elasticities with respect to price and inventory are calculated for every store-title-month observation and then averages are taken within each zipcode region - month and then across regions.

Table 7: Demand Results: Second Stage Regressions

	OLS Coefft (S.E.)	IV 1 Coefft (S.E.)	IV 2 Coefft (S.E.)
Release date:			
Quarter 2	-0.008 (0.004)	-0.012 (0.004)	-0.015 (0.004)
Quarter 3	-0.119 (0.004)	-0.120 (0.004)	-0.124 (0.004)
Quarter 4	-0.029 (0.004)	-0.040 (0.004)	-0.042 (0.004)
Box Office:			
B	-0.755 (0.042)	-0.685 (0.043)	-0.724 (0.045)
C	-1.363 (0.039)	-1.278 (0.041)	-1.317 (0.042)
Genre:			
Child/Family	-0.787 (0.051)	-0.600 (0.052)	-0.638 (0.051)
Comedy	0.172 (0.046)	0.265 (0.047)	0.250 (0.046)
Drama	-0.135 (0.023)	-0.069 (0.023)	-0.096 (0.023)
Horror/Suspense	-0.021 (0.031)	-0.021 (0.033)	-0.036 (0.034)
Romance	-0.821 (0.045)	-0.661 (0.046)	-0.688 (0.046)
Science Fiction	-0.721 (0.053)	-0.549 (0.056)	-0.530 (0.056)
Rating:			
PG	-0.021 (0.032)	0.027 (0.032)	0.044 (0.032)
PG13	0.052 (0.053)	0.171 (0.054)	0.198 (0.053)
R, NC17, NR	-0.006 (0.059)	0.115 (0.060)	0.124 (0.060)
Market characteristics:			
Median age	0.0073 (0.0006)	0.0084 (0.0006)	0.0067 (0.0006)
Median income	-0.0113 (0.0002)	-0.0114 (0.0002)	-0.0118 (0.0002)
Number of households	-0.0001 (3.4E-7)	-0.0001 (3.5E-7)	-0.0001 (3.6E-7)
Percent white	-0.0295 (0.0004)	-0.0289 (0.0004)	-0.0306 (0.0005)
Percent black	-0.0240 (0.0004)	-0.0233 (0.0004)	-0.0249 (0.0004)
Percent female	0.0163 (0.0020)	0.0128 (0.0020)	0.0209 (0.0021)
Percent single mother with kids	0.0063 (0.0013)	0.0064 (0.0013)	0.0079 (0.0014)
Percent single father with kids	-0.2546 (0.0037)	-0.2509 (0.0039)	-0.2526 (0.0040)
Percent single male	-0.3782 (0.0076)	-0.3797 (0.0078)	-0.3753 (0.0080)
Percent single female	-0.0382 (0.0043)	-0.0410 (0.0044)	-0.0390 (0.0045)
Percent married with kids	0.0196 (0.0005)	0.0194 (0.0006)	0.0227 (0.0006)
Percent with Bachelor's	-0.0115 (0.0003)	-0.0116 (0.0003)	-0.0110 (0.0003)
Blockbuster in market	0.6703 (0.0035)	0.6707 (0.0036)	0.6743 (0.0037)
Percent rural	0.0006 (0.0002)	0.0004 (0.0002)	0.0008 (0.0002)
Percent suburban	-0.0008 (0.0001)	-0.0009 (0.0001)	-0.0008 (0.0001)
N	407,006	407,006	407,006
R^2	0.43	0.40	0.40

Notes: Regression of estimated quality (including title fixed effect-store characteristic interactions, store fixed effects and all decay rate interactions) from nested logit on title and store characteristics. IV1 instruments for the within-group share only. IV2 instruments for both within-group share and inventory. Omitted category for Box is A; for Genre is Action/Adventure; for Rating is G. All specifications include interactions between title and store characteristics and between month fixed effects and title characteristics.

Table 8: Inequalities Analysis Results

	Coefft (S.E.)	95% Conservative CI
Per Tape:		
Constant	-11.3**	[-13.0, -11.3]
Box B title	12.9**	[12.5, 14.4]
Box C title	-0.46*	[-1.85, 0.09]

Notes: Results of inequalities methodology to estimate stores' costs of holding inventory. Coefficients represent predicted costs to the store per tape. "Box B title" and "Box C title" are indicators for titles in Box Office categories B and C: those with theatrical box office revenues $15-40 million and under $15 million respectively. **: significant at p=0.05; *: significant at p=0.10.

Table 9: Counterfactual Analysis Results

	1: Observed Contracts	2: Model predictions No delay		3: Model predictions No leverage	3: Model predictions With leverage
	$	$	% change from 1	% change from 2	% change from 2
Mean store profits	8,626	8,707	1.33%	-4.67%	-4.67%
Total FLF distributor profits	106,647	105,974	-0.64%	-29.53%	-29.53%
Total other distributor profits	1,336,466	1,348,916	0.93%	0.57%	0.59%
Mean consumer surplus per store	712	717	0.29%	-0.84%	-0.84%
Mean total welfare per store	26,725	26,952	0.95%	-1.87%	-1.85%
Mean titles taken per store					
FLF distributor: dropped / LP / RS			0 / 0 / 0	1.8 / 0.3 / 2.9	1.8 / 0.3 / 2.9
Other distributors: added on LP / RS			0.2 / 1.9	0 / 0	0.012 / 0.024

Notes: Column 1 reports model predictions under the observed store contract choices. Column 2 reports model predictions under predicted store contract choices with no delay in FLF introduction. Column 3 presents results of counterfactual analyses. Introduction of FLF contracts was delayed by three months for a single distributor. That distributor released 5 titles in the relevant window. 56 titles were released by other distributors in the same period. 83 stores are considered in each counterfactual.