The Effects of Full-Line Forcing Contracts:

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

March 7, 2008
Preliminary and Incomplete - Please do not Cite

1 Introduction

We consider the effects of full-line forcing or bundling contracts in the video rental industry. Studios offer their titles to video rental stores under four different contract types: linear pricing (LP), revenue sharing (RS), full-line forcing (FLF) and sell-through pricing (STP). Under linear pricing contracts the store pays a fixed price per tape, usually between $65 and $70. Under revenue sharing the upfront fee is much lower (around $8-$10 per tape) but the store also pays a fraction (in the region of 55%) of the rental revenues to the studio. The full-line forcing contract involves profit sharing under slightly better terms than the RS contracts but the store is also required to buy all the titles produced by the studio in the time period of the contract (usually 12 months). Finally, studios sell some titles on sell-through price terms: all buyers, including video rental stores, can purchase them for around $20-$25 per tape. There is no contract choice for these titles.

We discuss the intuition that in the absence of full-line forcing contracts, stores have an incentive to choose LP contracts for the titles for which they expect to have high demand, and RS contracts for other titles. This generates an adverse selection effect: the revenue sharing contract mitigates the inefficiencies generated by double marginalization, but does so only for low-demand titles, for which these inefficiencies are relatively small. The introduction of full-line forcing contracts may have three effects on the industry. First, since the store is required to take all the studio's titles on the same profit-sharing contract, some titles will be pulled out of LP contracts. This increases efficiency by moving titles from linear pricing to revenue-sharing contracts, increasing the use of two-part tariffs and therefore reducing the double marginalization problem and increasing the total profits to be divided between the studio and the store. We refer to this as the "efficiency effect" of the FLF contract. Second, if a store chooses a FLF contract when it would otherwise not have taken all of the studio's titles this will increase the number of the studio's titles which are available to the consumer in total. We call this the "market coverage effect". Finally, and conversely, the store may compensate for the requirement to take all of this studio's titles by dropping some titles produced by other studios, particularly if inventory holding costs are high. This is the "leverage effect" of bundling: the studio may offer a bundling contract for exactly this reason. 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

*Department of Economics, Harvard University, Cambridge, MA 02138 .Email: jho@fas.harvard.edu
†Department of Economics, Columbia University, New York, NY 10027. Email: katherine.ho@columbia.edu
‡Department of Economics, Harvard University, Cambridge, MA 02138. Email: mortimer@fas.harvard.edu

of the number of titles offered to consumers, the mix of studios producing those titles, prices and store and studio profits if full-line forcing contracts were not available. Our initial reduced form analyses show that the overall patterns in the data are consistent with the first two effects but that the leverage effect may be small. Selection issues - in particular those caused by stores choosing which contracts to accept - mean that a structural model is needed to analyze the market fully.

We model consumer demand for titles using a flexible nested logit framework which takes advantage of our large 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 studios. 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 and also the potential effect of a retailer's decision to take a full-line forcing or revenue sharing contract on future rentals of other titles, which as discussed below is not fully accounted for in our demand model. We estimate this cost using a method of moments methodology with inequalities, following the approach developed in Pakes, Porter, Ho and Ishii (2006). [We currently estimate a single cost that is common to all store-title pairs. We plan to extend this by estimating a reduced-form function for the cost including store, studio and market characteristics as independent variables.] Finally we [will] perform counterfactuals to investigate what the market would look like without full-line forcing contracts.

Our results suggest that...

This analysis does not focus on the question addressed in the theoretical literature on bundling: the reason why upstream firms might choose to offer bundling contracts. Our results are, however, informative on this question. There are three potential explanations in our context. The first is the efficiency issue discussed above: if titles are pulled from LP contracts into FLF contracts this increases efficiency by reducing the double marginalization problem, increasing the total profits to be divided between studio and store. The theoretical literature does not consider this possibility. The second is the leverage theory: bundling or tying may be used to "leverage" market power in one market to exclude competitors in another market. This theory has been discussed in numerous theoretical papers including, for example, Whinston (1990), Choi and Stefanadis (2001), Carlton and Waldman (2002) and Nalebuff (2004). The third explanation is price discrimination. If the preferences of the downstream firms (rental stores in our application) for each good are negatively correlated, then the upstream firm (studio) can profit by bundling goods together. Adams and Yellen (1976) provide the first formal model of price discrimination through tying, which was later generalized by McAfee, McMillian and Whinston (1989) and Salinger (1995) among others. Our demand model is sufficiently rich to accommodate this effect: a negative correlation in preferences across stores could be generated if the consumer population differs across markets and if preferences for particular titles vary across demographic groups. The demand model incorporates these effects. However, our data,which cover only 7 studios that offer 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.

There are to our knowledge very few previous empirical bundling papers. The first is by Chu, Leslie and Sorensen (2007). It studies bundling of tickets sold to consumers by a theatre company that produces a season of 8 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. There is also a small literature which 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). No previous authors have estimated a structural model of bundling in a supply chain setting.

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 considers the theoretical effects of tying in the video rental market. We begin by summarizing some important institutional features of the market. Then we discuss the theoretical implications of tying for efficiency in this industry.

2.1 The video rental market

The video rental industry has two primary tiers[1]. Studios make and distribute movies which are acquired by video rental stores who offer them for rental and sale to consumers. The studios use three different contractual forms for titles targeted to the rental market in their dealings with the rental stores. The first is linear pricing. Studios offer a title to a store for a fixed price 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. Studios enter into an agreement with the rental store to share the revenue generated by a title, in return for charging a reduced up-front fee. In the typical revenue sharing agreement the studio charges a upfront fee of \$3 to \$8 per tape and receives in the order of 55% of the rental revenue. The inventory decision of the rental store is often constrained by both maximum and (often binding) minimum quantity restrictions. Revenue sharing and linear pricing contracts both operate on a per title basis. That is, for each individual title, the rental store is free to choose both whether to purchase the title and which form of contract to take.

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 studio during the period of the agreement (typically 12 months) and to take them all under the same contract type[2]. In many other respects output programs resemble revenue sharing agreements. For each title, the studio receives an upfront fee per tape and a share of the revenues, both of which are usually lower than the revenue sharing terms. The quantity taken by the retailer is again restricted to be within a range, where the lower bound (on average 11 tapes per title) is again frequently binding.

Antitrust laws (Section 2 of the Clayton Act as amended by the 1936 Robinson-Patman Act) prevent distributors from offering different prices to different buyers for exactly the same product. They do not, however, apply to quantity restrictions or quantity discounts although copyright laws

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

[2] Some exceptions apply: titles released by the studio on "sell-through pricing" terms are exempt, and several studios 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 studio releases more than 3 titles in a month (a rare event), the retailer is only obligated to accept three of them. Finally, full-line forcing contracts also typically include opt-out clauses for movies with 'objectionable' content.

which permit stores to freely resell tapes purchased from studios under LP contracts put some bounds on this second degree price discrimination. In our application the data are consistent with a model where the studio must offer the same contractual prices (wholesale prices, upfront fees and revenue division) to all rental stores. The maximum and minimum quantity requirements for revenue-sharing contracts can and do vary with the box office of the movie and the size of the store.

In addition to setting contractual terms the distributor can in theory choose which of the contractual forms to offer. In particular one might expect the studio to choose not to offer linear pricing contracts since these are the least flexible contractual arrangements. In reality, however, full-line forcing contracts were not introduced until the middle of our dataset, in February 1999. Both these and revenue-sharing agreements require extensive computer monitoring of millions of transactions; only about half of the stores in the industry had the technology to adopt these contracts by 1998[3]. Thus studios were unlikely to choose not to offer linear pricing contracts during the period of our data since by doing so they would substantially reduce their target market[4]. This implies that rental stores can discipline the studios by opting to take linear-pricing terms when revenue-sharing splits are not satisfactory. The empirical evidence suggests that linear-pricing terms continue to be offered to all firms even when revenue-sharing terms 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 studio sells these movies to all buyers, including video rental stores, for quite low prices, often around $20-25 per tape[5]. This is much lower than the wholesale price under linear pricing contracts. 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 contract choice.

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

2.2 Empirical Implications of Theory on Tying

If demand is independent across titles, if only revenue sharing and linear pricing contracts are available, and if the same LP and RS terms are offered to all stores for a particular title, then retailers will choose linear pricing terms when expected demand for the title is above a certain cutoff value and revenue sharing terms for lower-quality titles[6,7]. Thus the efficiency loss from double marginalization may not be mitigated for high-value titles, for which the loss is relatively large. This is the source of the "efficiency effect" of a full-line forcing contract: since the contract requires

[3] Our dataset includes only stores that have the technology to do revenue sharing contracts.

[4] They were prevented from offering these 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 linear-pricing terms unless all retailers are excluded from these terms.

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

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

[7] We note in Section 5 that retailers take into account the contract-specific maximum and minimum quantity constraints set by studios, as well as the terms of the tariffs, when choosing between contract types. Since RS contracts have higher minimum quantity restrictions than LP contracts, and since a larger number of tapes per title on the shelf is likely to increase consumer awareness of the title, if stores prefer to advertise higher-quality titles then this may prompt them to take better titles on RS terms. This effect would go in the opposite direction from the effect of the type of tariff. However, the reduced form analysis in Section 4.1 demonstrates that the effect of the type of tariff dominates.

the store to take all of the relevant studio's titles under the same two-part tariff, valuable titles are pulled out of linear pricing contracts, which may significantly reduce the double marginalization problem[8].

We expect to observe two other welfare effects of introducing full-line forcing contracts. First, if the store previously took only a subset of the studio'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. (It may also be consistent with the price discrimination motive for bundling by studios.) Conversely, this effect together with the non-zero cost of holding inventory may prompt the store to drop other studios' titles: this is the leverage effect and is likely welfare-reducing since it reduces inter-studio competition.

The relative magnitudes of these three effects will depend on the mean and variance of demand for the titles produced by different studios and the extent of complementarities between them and also on stores' inventory holding costs. The aggregate welfare effect of full-line forcing 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 and monitors these contracts to facilitate payments between retailers and studios. The complete dataset combines information from previous studies (Mortimer 2007 and 2008) with additional information from Rentrak on full-line forcing 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 1025 titles released during these months.

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 store closure[9]. 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.

[8] The double marginalization problem would be negligible if retailers' marginal costs were very low. However, since each tape can be used only for a finite number of rentals before it breaks, and since inter-store competition implies a need to purchase multiple tapes in order to service demand quickly, each additional rental generates additional costs from purchasing new tapes. That is, we think of a tape as a "box containing τ_{jm} rentals", so stores' marginal costs include $\frac{1}{\tau}*$ the cost per tape, a non-negligible amount, in addition to the opportunity cost of holding inventory which we discuss below.

[9] For 1116 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.

For each title we observe a studio identifier (but not the actual studio name), its month of release to video, genre (such as Comedy or Action/Adventure), and MPAA rating (such as R or PG). 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 revenue-sharing and full-line forcing contracts offered to retailers for each title, and retail prices under linear-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. They provide information on the number of rentals per tape, on total weekly revenues per tape and on 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 is 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 dataset. We are left with 7,189 stores, 963 titles (201 in the A box-office category, 188 B titles and 574 C titles) and 59 studios in the dataset.

Full-line forcing contracts were first introduced in February 1999 of the dataset. 7 out of 59 studios offer a full-line forcing contract at some point in our panel. The average number of titles released per year by these 7 "FLF studios" was 10.5. The 52 "non-FLF" studios released on average 7.2 titles per year. Of the 7,189 retail stores in the clean dataset, 7,107 participated in at least one linear pricing contract during the period of the analysis, 6,687 participated in at least one revenue-sharing contract and 4,896 participated in at least one full-line forcing contract. On average stores take 42% of the titles released per studio-year. The proportion is higher for FLF studios (69%, or 43% if we include only stores that have no FLF contracts with the relevant studio) than for other studios (38%). That is, FLF studios both release more titles per year and also have higher take-up rates than other studios.

Stores are categorized into ten tier sizes, 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 12.2 in tier 1 to 15.3 in tier 6. It then falls in each subsequent tier to a low of 10.3 in tier 10. The average number of tapes taken per month increases with every store tier, from 58.7 in tier 1 to 360.1 in tier 10. The average number of titles taken per month differs very little between stores that take FLF contracts and those that do not. However, "FLF stores" take more tapes per month than "non-FLF stores" (165.6 with standard deviation 207.5 compared to 105.8, standard deviation 163.0).

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 price for linear pricing contracts is \$66.82, compared to an average upfront fee of \$8.48 for revenue sharing contracts, \$3.60 for full-line forcing contracts and a price of \$15.17 for sell-through price contracts. Retailers on average keep 46% of revenues under revenue sharing contracts and 59% of revenues under full-line forcing 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 and 22 respectively. Average month 1 rentals are highest under revenue sharing contracts but the decay rate is also greatest for these titles; by month 3 linear priced titles have higher demand and this remains true in months 4 and 5. Average rental prices differ very little across contract types. Not surprisingly, retailers extract the largest number of rentals per tape for titles purchased under linear pricing contracts. Average inventory levels are highest for titles purchased under sell-through pricing and revenue sharing contracts and lowest for those under linear pricing contracts.

Tables 2 and 3 summarize the numbers of titles offered by studios, and taken by stores, under different contract types. The majority of titles in our data were offered under linear pricing contracts; approximately 70% were also offered under revenue sharing contracts. No full-line forcing contracts were offered in the first year of our data; a total of 10 were offered in year 2, 18 in year 3 and 39 in year 4. Table 3 shows 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 adopted 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 choose to 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. 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. We ask what percent of stores in the lowest quintile of % of titles adopted under LP contracts are in store tiers 1-3. We then normalize by the percent of all stores that are in those tiers. The result (a figure of 1.35) indicates that small stores are over-represented in the first quintile of LP contracts. Overall, small stores are over-represented in the first and second quintiles of LP contracts, the first 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 fifth quintiles of LP and FLF contracts and in the first and second quintiles of RS 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. Full-line forcing contracts are not necessarily predicted to be correlated with low demand. 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 are similar to those choosing LP. We expect a similar pattern here. We regress store revenues on an indicator for the adoption of a revenue-sharing contract, an indicator for the adoption of a full-line forcing contract and title fixed-effects at the store-title level. Consistent with our prediction, we find that revenues for Box Office group A titles are approximately \$109 lower under revenue-sharing than under linear pricing contracts (standard error of 1.33) and that revenues under full-line forcing are not significantly different from those under linear pricing (coefficient of -10.15, standard error 5.47).[10] The relationship between RS and LP revenues is similar for Box Office group B and C titles.

Next we investigate in more detail which types of stores choose full-line forcing contracts. A logit regression of a dummy for participation in these contracts on observable store characteristics indicates that larger stores and those in suburban areas are more likely to adopt full-line forcing contracts. Stores facing more competition from a Blockbuster Video are more likely to adopt full-line forcing.

Finally, we would like to investigate whether the introduction of revenue sharing and full-line forcing contracts had a positive effect on market coverage. First the statistic noted earlier, that stores on average take 43% of titles released by studios that offer FLF contracts at some point, excluding FLF contracts themselves, is consistent with a reasonably large effect. Second, we consider a number of summary statistics separately for different zip code groups (since the dataset is too large for us to run some of the regressions using all the data together). For stores in the first zip code category[11], the average number of titles taken per month from studios from which the store takes a FLF contract in some other month is 0.59. The average increases to 1.32 titles per month in FLF months, suggesting a positive market coverage effect. We regress the number of titles taken per studio-month, the number of tapes per title and the number of transactions per title on an indicator for active FLF contracts for the relevant store-studio pair and store fixed effects. We include only store-studio pairs for which a FLF contract exists at some point in our panel. We are therefore looking within-store and asking whether taking a FLF contract from a particular studio is correlated with inventory choices, and rental activity, specific to that studio. The coefficient on FLF activity in the regression considering the number of titles taken per studio-month is positive and significant (coefficient 0.76, standard error 0.02, from a mean in the data of 0.70). That in the regression considering the number of tapes per title is 1.01 (standard error 0.24 from a mean of 12.5). This implies a positive market expansion effect in terms of both the number of titles taken and the number of tapes per title. However, the equivalent coefficient in the transactions regression is negative and significant (coefficient -63.0, standard error 4.18 from a mean of 231). The new titles taken under the FLF contract generate fewer rentals than other titles. The results are similar across zip code categories.

[10] The results imply that the minimum quantity restrictions for RS contracts, which might prompt stores to choose RS terms for high quality titles and might also, by increasing consumer awareness of these particular titles, themselves increase demand for RS titles, are outweighed by the effect of the difference in tariffs between LP and RS contracts.

[11] This zip code category contains zip codes from 20000 to 24999. It includes areas on the East Coast such as Washington DC and parts of Virginia.

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 studios. We might expect this effect to generate a negative correlation between the adoption of full-line forcing contracts by a retailer and the orders (or rentals) of products from other, non-bundling studios. 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 may alter the intuition.

First we note that, for stores in the first zip code category, the average number of titles taken per month from studios from which the store never takes a FLF contract, in months where it has no active FLF contracts, is 0.281. The equivalent number for months in which the store has some FLF contracts (with other studios) is 0.224. The raw data therefore indicate a very small leverage effect in terms of the number of titles taken. We then regress the number of titles taken per studio-month, the number of tapes per title and the number of transactions per title on an indicator for active FLF contracts in this month with some other studio. We include store fixed effects and exclude from the regression studios with which the store ever has a FLF contract. The coefficient on the number of titles taken per studio-month is negative and significant but very small (coefficient of -0.064, standard error 0.003 compared to a mean in the data of 0.24, standard deviation 0.75). The coefficient on the number of tapes per title is 0.21 (standard error 0.10 from a mean in the data of 9.7). That on transactions per title is -21.9 (standard error 1.74, mean in the data of 183.2). Overall, then, the simplest reduced form regressions imply that there may be a leverage effect in terms of the number of titles taken per studio-month, but this is small and the number of tapes taken per title may even increase. The number of transactions per title from other studios falls somewhat, presumably due to a substitution effect when additional titles are added from the FLF studio. The results from other zip codes are consistent with this summary.

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. A structural model is needed to correct for the selection problem arising from stores choosing which contract types to take. For example, large stores are more likely than others to select FLF contracts. This issue makes it difficult to predict the effect of introducing new FLF contracts from the reduced form results alone.

The modeling approach we propose 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, changing the number of tapes on the shelf and changing the price per rental. The second step is to use moment inequalities to infer the store's cost of holding inventory according to the overall size of the store. The use of moment inequalities over contract choice means that the issue of selection into different contractual forms is explicitly modelled and thus should not hinder identification. 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 studios and stores, and to assess the competitive impacts of the partial adoption of FLF contracts across the industry.

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 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[12]. 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 by potentially 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. 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 two other potential effects of an increase in inventory for a particular title X. First, if some other title Y is a close substitute for X released in the same month, then its rentals are likely to fall (and in fact the store may choose not to take it at all). This could imply an overall reduction in retailer profits for this pair of titles, for example if title Y would have been taken on LP or sell-through pricing terms but X was taken on RS or FLF terms so that the proportion of the rental revenues captured by the store was higher for Y than for X. Our model will incorporate this effect since as discussed below the demand system models consumers' choices between the titles available in each particular month, taking inventory levels into account, and our supply model allows retailers to consider specific contract terms when choosing their strategies.

Finally there is a cross-month effect that is not accounted for in our demand model. If consumer preferences are correlated across months this will affect the impact of a change in X's inventory level on retailer profits. For example, suppose that titles X and Y are both romantic comedies. Title Y is released six months after title X's release. Each household contains two decision-makers, one of whom prefers romantic comedies and the other prefers action movies, implying a negative cross-time correlation in preferences so that renting title X in its release month implies renting an action movie rather than title Y in a later month. Titles X and Y can therefore be thought of as substitutes even though they are released in different months. Then if the store takes title X on a FLF or RS contract, implying a high inventory level, it is likely to have high rentals of title X and low rentals of title Y, reducing the retailer's revenues from Y compared to an alternative where X is taken on a LP contract. If rentals of Y are more profitable to the retailer on the margin than rentals of X (for example because Y is taken on a LP contract) and if the potential action movie rentals in the two months are equally profitable to the retailer and less profitable than title Y, then this implies a reduction in retailer profits. Our demand model will not capture this effect: it incorporates cross-month effects within-title and within-month effects across titles but does not allow a change in rentals of one title to affect later rentals of a different title. The average impact of this effect on retailer profits will be included in the cost of taking an extra tape estimated in the

[12] The relatively high base cost per tape under LP will also tend to reduce the number of tapes taken per title.

[13] A second mechanism through with 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.

inequalities analysis. It may imply a non-trivial estimated cost of taking a RS or FLF contract[14,15]. The intuition for how the average effect is identified is discussed in more detail in Section 7.

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.[16]

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 which are unrelated to true demand; allowing consumers to substitute across weeks within each month mitigates this problem. Unfortunately 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 (2007) 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[17]. Our methodology is as follows. For any title released in month one, we summarize over weeks 1 through 4 to generate month 1 demand[18]. 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

[14] The sign of this "cost" would change if X and Y were complements rather than substitutes so that an increase in rentals of X led more consumers to rent Y. It could change if increasing the inventory of X increased the size of the market rather than just causing substitution between titles. It also depends on the terms on which title Y is taken. The effect is zero if Y is taken on the same terms as X. It also cancels out if Y and the two action movies are equally profitable to the retailer, since then modeling substitution across titles within month 1 has the same result as modeling the true cross-time substitution. Since only RS and FLF contracts have minimum quantity restrictions, and since the modal contract type is a LP contract, the case where X is taken on RS or FLF terms and Y taken on LP terms is the most likely. The mix of contract terms between Y and the action titles, as well as the other open questions noted here, are empirical questions. The net effect will be estimated in the inequalities analysis.

[15] There is another cross-month effect that goes in the opposite direction. If X and Y are released in the same month and the store chooses to stock title Y despite the reduction in demand in its release month, then consumers who rented title X in month 1 would be more likely to rent Y in the following month than they would have been before X's inventory increased. That is, because of the durable goods aspect of demand, title Y's decay rate is likely to be affected by an increase in X's inventory. Our demand model contains an average decay rate that varies by type of title but does not identify these decay rate changes. If Y substitutes for titles other than X in month 2 then in expectation it will not affect retailer profits since on average the two titles have the same contract terms and similar prices. However, if Y substitutes for X this could increase retailer profits. We expect this to be a second order effect since substitution from X would occur only for consumers who would have rented X twice.

[16] We assume that, if the store switches to taking a FLF contract, it places all the studio's tapes on the shelf rather than storing some in a back room. If this assumption is incorrect our inequalities analysis will estimate a very low cost of holding inventory under FLF contracts. In addition our demand predictions for FLF titles will be biased up (since they assume that all tapes are present on the shelf) and that for other titles will be biased down. The overall measurement error in store returns will be captured as part of the cost estimated by the inequalities analysis.

[17] 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.

[18] We replace missing or negative values for revenues or transactions with zeros. Average prices are constructed by dividing monthly revenues by the number of transactions in the month.

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[19]. We construct prices at the monthly level by dividing monthly revenues by monthly transactions. We observe very little variation in price within box office class-contract type groups, either for studios or for stores. We will assume throughout our model that price is fixed within these cells[20].

We define the 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 have weak substitution with current releases. Only these titles are included in the analysis for the relevant month[21,22]. We group titles according to genre (action/adventure, children/family, comedy, drama, horror, romance and sci fi) and box office class (A versus B versus C). Industry discussion indicate that A titles are likely to be more heavily advertised and displayed more prominently than B or C titles.

We estimate a nested logit model of demand with nests defined as genre/box office class groups[23]. 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, 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 $[\zeta_{igt} + (1-\sigma)\varepsilon_{ijmt}]$ 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:

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

where δ_j is a title fixed effect, η_m is a store fixed effect, θ_t is a month fixed effect and p_{jmt} is the average price per rental of the tape in month t. 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

[19] 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, which will bias down the demand estimates for those titles. Fortunately, there appears to be no correlation between this and studios, 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. [Check: dropping month 5+ observations didn't affect our demand estimates?]

[20] Industry conversations indicate that retailers view the price of a new release as being essentially fixed, but that different contract types imply different numbers of tapes. A higher inventory implies an increased length of rental (and lower price per day) and therefore fewer late fees. We capture this effect by using average revenues per rental as our price variable. Our inequalities methodology assumes that stores essentially follow industry rules for average prices within contract type and box office class.

[21] Our dataset includes titles released between months 1 and 52 of our panel. Titles released between months 1-4 and between months 49-52 compete with those released in months -3 to 0 and 53-56 respectively, which we do not observe. We therefore exclude months 1-4 and 49 and above from the final demand and inequalities analysis, ensuring that we include only months for which we observe the full choice set.

[22] We choose to pool the data 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. The cost of our approach, which uses a nested logit specification, is that we estimate a single σ parameter rather than letting this vary across months.

[23] 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 class are the main sources of differentiation between titles this seems a natural structure for modeling heterogeneity in consumer preferences.

the demographics of the store's market. The variables are the per cent white, the per cent single and the per cent with children. We therefore permit each store to predict the demand for a particular title based on the demographics of local consumers[24].

The decay rates θ_t capture 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[25]. 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, rating (G, PG, PG13 and R+) and double and triple interactions of these three groups of variables)[26]. This implies constraining the decay rate to be the same for all titles in a particular box office class-genre-rating cell[27]. 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:

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

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 is equal to 4 movie rentals per household per month in the zip code of the store.

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 if it is the only one we observe 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. The same assumption will be used in the inequalities analysis: each store evaluates its alternative contracts and portfolios without predicting potential responses by other same-market stores.

We could alternatively have included all observed stores in each market in the estimated demand system and 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. However, this methodology would have

[24] 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 "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.

[25] 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.

[26] 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.

[27] 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.

been significantly more complicated to implement. We justify our methodology by assuming that store competition occurs through location and an average price measure. Once a consumer enters a particular store he is very unlikely to shop around for different stores on that particular day[28]. The relevant dimension of competition, particularly since we are considering bundling, is that across movie studios within a store rather than that across stores. We model the former carefully but do not go into details on the latter.

Of course this simplification may lead to biased results. This seems most likely if, in reality, stores choose particular contract types in order to steal business from their competitors. However the fact that retailers in the same market choose the same contracts more than 90% of the time makes this argument less convincing[29]. The largest concern is with Blockbuster, which has FLF contracts for almost all titles and therefore frequently has a different contract type from 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). We would like to include a Blockbuster dummy variable to account for the fact that the equal allocation of potential customers across stores is not quite right in this case; this variable is absorbed into the store fixed effects[30].

Three variables in the demand model are likely to be endogenous: the price variable, inventory and $s_{jmt/gmt}$, the share of the title within its group. Since the demand model includes both store and title fixed effects we are concerned about endogeneity only through unobservables that change over time and affect changes in prices, in inventory and in the $s_{jmt/gmt}$ term. We instrument for inventory using the average inventory for titles in the same store-box office-contract type group that were released in other months. The assumptions needed to make this a valid instrument are that overall inventory for a particular group of movies is related to store shelf space (i.e. to store-specific costs) and that other titles in the same group that were released at different times have demand shocks that are not correlated to those of this title. The instrument will probably be invalid if demand shocks are correlated over time. We tried using other instruments such as contract type but encountered problems because the numerous fixed effects in the model leave very little variation to be used for instrumenting.

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[31]. 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 $\ln(s_{jmt/gmt})$: it is a valid instrument under the assumption that demand shocks, which might affect the share variable, are not correlated across

[28] Over time, if one store improves its offering it will gain market share. However, the fact that we very rarely observe two or more stores per market means we do not have enough data to identify this effect.

[29] That is, if one store switches to FLF its competitors will also switch, so it cannot assume that the new contract will differentiate it from the other stores in the area. The potential for there to be multiple equilibria may cause problems here, since other equilibria besides the symmetric outcome may be possible in some regions of the parameter space. Like many previous authors, we do not attempt to address the multiple equilibria problem in this paper.

[30] One problem still remains. If a store chooses FLF in order to steal business from Blockbuster then the coefficient on the dummy variable should change in the counterfactuals discussed below. We cannot account for this given the available data.

[31] 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.

markets.

We tried numerous instruments for price, including measures of variable costs and average prices of other similar titles. None of the instruments were successful. The issue is that, after including store, title and month fixed effects, the only unobservable we need to instrument for is at the store-title level. There is little variation in price at this level. 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[32].

Some titles are observed to have zero revenues (i.e. zero rentals) in one or more months for a particular store. We drop zero revenue observations where positive revenues have not yet been observed since it seems likely that these zeros are caused by the store introducing the title after its release date. Zero revenue title-months for which no positive revenue months are seen later in the panel are also excluded since the store may have permanently removed these titles from its shelves. Finally, we also drop the 1.5% of store-title-month observations that have zero revenues in that month but positive revenues in both earlier and later months: we interpret these as months when the title was not on the shelf or when transactions were not recorded for some reason.

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[33].

6.2 Demand Results

We report results for the first zip code category in Table 5[34]. 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 class, genre and rating and interactions between these); the results are not reported due to space constraints. 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. This good fit with the data is particularly useful since our supply side estimation will stay within-sample in terms of titles and stores, allowing

[32] We also estimate a demand equation using share of revenues, rather than share of transactions, as the dependent variable. We use exactly the same specification as in our primary model, but define each title's market share in a particular store-month as the revenue to the store from this title-month divided by total potential revenues (the average price over all titles for the relevant store-box office group multiplied by the same market size used in the main model). We use this approach rather than defining the denominator as the total revenues in the store-month in order to allow consumers to choose the outside good of not renting a title. This method implies an assumption that stores do not cut prices in order to attract new customers, which seems unlikely given that observed prices differ very little across titles. The estimated coefficients (on all variables other than price, which of course is excluded from the specification) are very similar to those from the main model. We could in future repeat our supply side analysis using the results from this specification. This would act as a robustness test of our assumption that there is enough price variation left in the data to estimate a reliable price coefficient that can be used to predict store revenues under counterfactual contract terms. (Of course this robustness test would not permit us to analyze consumer welfare without making an assumption about price elasticities of demand.)

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

[34] As noted in Section 4.1, this zip code category contains zipcodes 20000 - 24999. It includes, for example, Washington DC and parts of VA.

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. The estimated decay rates are also 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 coefficients on inventory and the interactions between inventory and the decay rate. 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 first intuition.

We repeated the demand analysis for each of the 14 other zip code categories. The price coefficient was negative, as expected, for all but one of these: the problematic zip code category 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. There is some variation in results across the remaining 14 zip code categories. The (unweighted) average price coefficient is -0.027. Its standard deviation across zip code categories is 0.015; the minimum value is -0.051 and the maximum is -0.002. The average inventory coefficient is 0.017, standard deviation 0.005, minimum 0.009 and maximum 0.028. The R^2's in the regressions range from 0.757 to 0.792.

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 zip code category and each month since release. We then take an unweighted average over zip code categories to generate our summary data. The average elasticity with respect to price over all months since release is -0.141. The average elasticity with respect to inventory is 0.181. As for the main demand results, there is some variation across zip code categories. The standard deviation in price elasticities across zip code categories is 0.081; the minimum value is -0.260 and the maximum is -0.011. The standard deviation in inventory elasticities is just 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 Groups. These patterns are also displayed in Figure 1. The price elasticity is always negative and varies over time, for every Box Office group, according to a U-shaped curve: the elasticity increases slightly (becomes slightly more negative) every month from month 1 to month 4 and then becomes slightly less negative in month 5. The inventory elasticities 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 groups, and for each group the elasticity decreases over time since release. These results together 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: they want to rent high-buzz titles and do not care too much about small price differences across

titles. 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. They may be more influenced by price than other consumers. However, the variation in price elasticities across months and Box Office groups is fairly low. The limited price variation left in the data after including the fixed effects in the demand specification may relate mostly to issues such as late fees rather than the price paid upfront by consumers, implying that our elasticity estimates are not very informative about true consumer responses to variation in ticket price[35].

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 $\left[\hat{\delta}_j + \hat{\gamma}_j z_m + \hat{\eta}_m + \hat{\theta}_t + \hat{\beta}_t x_j\right]$. Results are reported for the first zip code category; those for other geographic areas 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 $\theta_t + \beta_t x_j$ term that was 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 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 only high-demand markets[36]. The R^2 on these regressions is only 0.4: even with 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.

[35] The increase in price elasticities in month 5 may indicate that consumers who rent a title many months after its release are looking for it specifically and do not view other titles as adequate substitutes. However, the pattern may also be due to our decision to aggregate demand from all later months into the month 5 observation. We could as a robustness test re-estimate demand including 6 or 7 separate months for each store-title pair and note whether this final-month increase in elasticity disappears.

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

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, and also the effect of an increase in inventory of title X on later rentals of other titles which is not incorporated in the demand model. This may imply a non-trivial estimated average cost of taking FLF or RS contracts.

We use the method of moments inequalities estimator developed in Pakes, Porter, Ho and Ishii (2006) to estimate inventory holding costs. That paper shows how to use the 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 contract types 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 will model 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 contract type 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 derive inequalities from every store's choice of contract for every title. We assume that the store may imperfectly predict the cost of holding inventory but that it has perfect foresight regarding the titles to be released to video in the following five-month period, consumer demand for these titles, and the contract type on which it will take each title (and therefore the number of tapes chosen per title). 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 have substantial amounts of information about each title several months in advance. Our assumption, noted above, that each store operates as a monopolist in its own 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.

The intuition for identifying the cross-time demand effects noted above is straightforward. The demand model does not account for the potential effect of an increase in one title's inventory on future rentals of other titles. The predicted retailer returns from switching to a contract which requires an increase in inventory of some titles (e.g. switching to RS or FLF contract) may therefore be biased: in the example discussed in Section 5 this prediction would be biased upwards. The difference between the model's predicted returns and the true expected returns that inform stores' choices will be captured in the estimated "cost" of holding additional tapes.

7.1 The Store Profit Equation

Our first step is to predict the total return to the store from its contracts with all studios 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:

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

where:

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

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 studio. We denote the return from title j under the three contract types as follows:

1. Under linear pricing the return for title j is $r_{jm}(.) = \left(\sum_{t=t_j}^{t_j+4} q_{tjm}p_{tjm}\right) - F_j c_{jm}$. 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 we write $r_{jm}(.) = y_j^{RS}\left(\sum_{t=t_j}^{t_j+4} q_{tjm}p_{tjm} - u_j^{RS}c_{jm}\right)$, 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 produced by the studio during the following twelve months. It receives slightly better terms than those under revenue sharing: $r_{jm}(.) = y_j^{FLF}\left(\sum_{t=t_j}^{t_j+4} q_{tjm}p_{tjm} - u_j^{FLF}c_{jm}\right)$.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 tape, τ_{jm}.[37] Additional constraints, in the form of minimum and maximum quantity restrictions on inventory purchases, are also set by the studio for RS and FLF contracts (at the store-title level). We denote these constraints as $\underline{c}_{jm}$ and $\bar{c}_{jm}$ respectively. Then 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:

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

[37] 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 maximum number of rentals per tape that the store is observed to provide for titles in the same box office class 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.

where time is measured in months. As before we consider the first 4 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 this function $r_{jm}(.)$, we can write the store's profit from its observed contracts as:

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

where J_s is the set of titles produced by studio s during the time period covered by our data and C is the inventory cost of holding each tape. When the store chooses not to stock a title, we assume it makes no contribution to the store's profit. η_m is a store fixed effect, k_{ms} is the contract type (a vector with one element per title for this store-studio 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[38]. The final term, ε_{ms}, is an unobservable such as store prediction error in the cost of holding inventory. [We plan to estimate a more detailed reduced form function for inventory costs, including independent variables at the store, studio and title level. This should be sufficiently detailed to capture any cross-store or cross-month variation in inventory holding costs or in the demand effects that are not included in the demand model. This analysis has not yet been completed; the remainder of the text considers a single estimated cost that is common to all stores.]

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[39].

7.2 The Inequality Estimator

Now consider the store's choice of contracts. Take as an example a title t' released by studio s' for which the store chose a LP contract. It could instead have chosen a RS contract; it could also have chosen, at the time when this title was released, to switch to a FLF contract with the studio for the following year. We assume that:

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

for the observed and any alternative contract choices, where π_m^{alt} is defined analagously to π_m^{obs}, but using a set of alternative contracts defined below. The expectation is taken conditional on I_m, the store's information set at the time when it makes its choice[40]. We therefore infer from the observed data that:

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

[38] As a robustness test we could choose alternative portfolios whose end-of-period contracts are the same as those observed in the data. $\rho(\tilde{c}_{ms}, k_{ms})$, like η_m, would then be precisely differenced out of our inequalities. However this would imply dropping a significant proportion of our data on FLF contracts.

[39] We plan to test the sensitivity of our results to variation in these cutoff levels.

[40] We consider only alternative contracts that were offered by the studio for the relevant title and only titles that are in release in our data for at least six months (except in the case of a switch to or from a FLF contract, where we include all titles released in the relevant year).

where $k^{LP}_{ms'}(t')$ indicates that the t'th element of $k_{ms'}$ is a LP contract. This equation implies the following inequality (assuming that title t' has zero demand by the end of the 4-year period in the data and therefore that the $\rho(.)$ term is differenced out[41]):

$$E\left(\Delta\pi_m^{s',t'}(.) \mid I_m\right) = E\{\sum_s[\sum_{j\in J_s}\{\Delta r_{jm}(k^{LP}_{ms'}(t'),k^{RS}_{ms'}(t'))-C*\Delta\tilde{c}_{jm}(k^{LP}_{ms'}(t'),k^{RS}_{ms'}(t'))\}+\Delta\varepsilon_{ms'}] \mid I_m\} \geq 0 \tag{13}$$

Here the difference function $\Delta\tilde{c}_{jm}(k^{LP}_{ms'}(t'),k^{RS}_{ms'}(t')) = \tilde{c}^{obs}_{jm} - \tilde{c}^{alt}_{jm} = \tilde{c}_{jm}(k^{LP}_{ms'}(t')) - \tilde{c}_{jm}(k^{RS}_{ms'}(t'))$, and similarly for $\Delta r(.)$. The returns from the observed and the alternative portfolios are calculated from the model (even though the observed portfolio return is observed) to ensure comparability in the counterfactual[42]. Each observed choice of contract implies a similar inequality.

We form inequalities for estimation by taking an expectation conditional on the instruments $z_{ms'}$, where s' is the studio whose contracts 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:

$$E(\Delta\pi_m^{t',s'}(.) \mid z_{ms'}) = E\left\{\sum_s\sum_{j\in J_s}\left(\Delta r_{jm}^{t',s'}(.) - C*\Delta\tilde{c}_{jm}^{t',s'}(.)\right) \mid z_{ms'}\right\} \geq 0. \tag{14}$$

The variables in $\pi_m(.)$ that will change across alternative contract types are the contract terms, the indicator functions I_{jm}, prices p_{jtm}, capacities c_{jm} and technology (rentals per tape) τ_{jm}. The terms of the potential contracts, $(F_j, y^{RS}, u^{RS}, y^{FLF}, u^{FLF})$, are calculated as the modal values over all stores for that contract type and title[43]. Our assumptions regarding prices and quantities are as follows. We note that prices vary only slightly between contract types 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 contract types. Instead we use the average price for each month for the particular box office class-store-contract type combination being considered[44].

Similarly, we do not formally model the firm's choice of c_{jm} and τ_{jm} for every title and contract type. We define the quantity $\tilde{q}_{tjm}$ as in equation (7). The first term, $Ms_{tjm}(.)$, represents consumer demand for the title in month t. We predict this using the estimated demand coefficients, the titles offered by the store under the relevant contract type, and the price and inventory choices for this contract type and title, defined as averages over other titles in the same store-box office class-contract type-month. The inventory level is also constrained by the maximum and minimum

[41] Note that when considering FLF contracts, if the contract has not expired by the end of our panel, our alternative contracts will return to FLF at the end of the data. This is not quite correct since the contract would then have longer left to run than does the observed contract. This will add noise to the estimates since the $\rho(.)$ terms may not exactly cancel out in this case.

[42] In fact we calculated 8 possible returns from observed contract choices, using every possible combination of observed/expected price, demand and maximum supply.

[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 studios. 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 either observed or alternative contracts.

[44] Our price measure is the average revenue per rental: it therefore incorporates the variation in rental periods (and therefore in price per day) and late fees that come from inventory variation across contract types. We assume that stores take this variation as given when choosing contracts. We do not have a clean measure of rental periods that can be compared across stores so do not have the data to model the choice of price per day or of rental period for each title.

quantity restrictions for that title as defined in equation (8)[45]. 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 highest τ observed for titles in the same store-box office class-contract type. The inventory level is the same value used as an input into expected demand[46,47].

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 fee, wholesale price and maximum/minimum quantity restrictions) are defined by the studio and perfectly observed by the store. We use the modal values for the relevant title-contract type across stores (in cases where regulations require the studio to charge all stores the same amount) or across same-tier stores (in other cases, i.e. for minimum/maximum quantity restrictions) in order to remove any special terms (such as quantity discounts) that are negotiated by a small number of stores. We predict expected prices using averages over titles within the same store-box office-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)[48]. 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, we follow a methodology consistent with that used in the demand model. Zero transactions at the beginning or end of the panel for a particular title are interpreted as being caused by the store choosing not to offer the particular title: we therefore set expected price and expected transactions equal to zero. Zero transactions in middle months are also excluded to ensure consistent comparisons between the observed and alternative contract types[49].

[45] 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.

[46] If a store takes a title under the FLF contract that it did not take before, we assume it purchases the minimum number of tapes required by the studio (rather than the average for that store-box-contract type). If the store-box-contract group is ever empty we again use the minimum number of tapes required by the studio 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 groups. 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 taking under LP. We predict price and τ_{jm} in the case where the store-box-contract group is empty using other contract types. Price for a title taken on a FLF contract that has never been taken before by the store is defined as the minimum of the average prices in the same store-box office group under LP, RS and STP contracts. (We use a minimum because titles newly taken under FLF have lower average quality than those observed to be taken by the store under 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 average prices 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 group for neighboring months.

[47] When considering the observed contract choice, our "observed demand" is calculated from the estimated demand coefficients together with the observed price and inventory levels. The "observed maximum number of rentals" is the maximum τ_{jm} observed for same-store-box-contract type titles multiplied by the observed inventory level. The expected values are calculated using exactly the same methodology as those for alternative contract types.

[48] We could have considered titles in the same store-genre-box office-contract type-month, consistent with the nest definitions in the demand model, but encountered problems with small sample sizes in some cases. The other expectation used in our model is an average across same-tier stores for the same title, used to define instruments in the demand model.

[49] This implies an assumption that the store does not change the amount of time each title is offered to consumers when it changes contract types.

The last step is to choose the alternative contracts that will define the inequalities used for estimation. We use the following framework. First we ignore FLF contracts and consider the store's choice between LP and RS contracts for each title[50]. We estimate inventory holding costs using just the inequalities defined by those choices. Then we add the inequalities generated by considering the store's choice between FLF and other contract types [this part is not yet completed].

Consider first the inequalities generated by the choice between LP and RS contract types. We define the store's alternative to the observed contract for any particular title t' from that studio-year to be switching from LP to RS or vice versa. For each t' we calculate $r_m^{altt'}(.) = \sum_s \sum_{j \in J_s} (r_{jm}^{altt'}(.))$, the store's total return when it switches the contract for title t' from LP to RS or vice versa, holding all other contracts and inventory levels fixed[51]. The calculation incorporates the revenue from all titles offered by all studios in all years in the data, since changing a single contract may affect demand for other-studio titles, even if these are offered in later months[52]. We generate the equation for estimation by converting the expectations in equation (14) into sample averages over stores[53]. We also average over alternative contracts t' in a particular studio-year before interacting with the instruments since the store may maximize profits at the studio rather than the title level. This implies the following equation for estimation:

$$\Delta\bar{\pi}_{ys'}^{noFLF} = \frac{1}{M}\sum_m \left((r_m^{obs} - C * \tilde{c}_m^{obs}) - \frac{1}{Q_{ys'}} \sum_{t' \in (s',y)} (r_m^{altt'} - C * \tilde{c}_m^{altt'}) \right) \otimes g(z_{ms'}) \geq 0 \qquad (15)$$

where y indexes years, s' indexes studios, $Q_{ys'}$ is the number of titles offered by studio s' in year y, $g(.)$ is any positive-valued function of the instruments, M is the number of stores in the data and r_m is the sum of r_{jm} over all titles. The total store-level inventory holding cost for a particular set of contracts is written as:

$$C * \tilde{c}_m = \sum_s \sum_{j \in J_s} C * \tilde{c}_{jm}(.).$$

We therefore have one moment per studio-year-instrument triple. We exclude studio-years where no RS contracts are offered and average within each year (before taking the store average) over studios that released fewer than five titles during our panel. This generates 40 studio-year moments per instrument. 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. Our results are set out in Section 7.3.

[50] We include titles where FLF was offered but exclude stores that chose the FLF option. We could alternatively exclude titles released by studios that offered FLF in the relevant year, since that would help avoid any selection effect created by endogenously-chosen stores in those years.

[51] If the alternative contract choice 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 title that is above the minimum quantity restriction of the contract; otherwise the alternative contract is defined as dropping the title.

[52] 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.

[53] We stack the inequalities for all zip code categories before taking the average over stores. We therefore estimate a single set of costs for all stores, taking into account all of the zip code category-specific demand estimates from Section 6.

Our methodology is slightly different in cases where the studio does offer a FLF contract in a particular year. First we note that, of the seven studios that offer FLF contracts during our data, six seem to have offered stores trial periods during which they could evaluate the profitability of FLF contracts. That is, many stores initially took just a few titles from the studio on FLF terms, then switched to other contract types before (possibly) switching back to FLF[54]. We define stores as taking FLF only if they adopt FLF after the trial period is over and if the total span on FLF is a full year (or lasts up to the end of our dataset, whichever is shorter). The start date is defined as the first month of FLF[55].

We define the inequalities for estimation, for stores that take FLF from a particular studio, by considering their choices for the duration of the FLF contract. The alternative is defined as taking all titles that were previously part of the FLF contract on RS contracts[56]. If the relevant store never takes a FLF contract from the relevant studio, we consider its choices for the 12 months after FLF was first offered. The alternative is taking all the studio's titles for that year on a FLF contract[57]. We calculate $r_m^{altys'}(.)$, this time defined as the store's predicted return when it changes all contracts from studio s' throughout year y. Then we form the following moments to add to those from non-FLF studio-years:

$$\Delta\bar{\pi}_{ys'}^{FLF} = \frac{1}{M}\sum_m \left[\left((r_m^{obs} - C_m^{tier}\tilde{c}_m^{obs}) - \left(r_m^{altys'} - C_m^{tier}\tilde{c}_m^{altys'}\right)\right) \otimes g(z_{ms'})\right] \geq 0. \tag{16}$$

This creates one additional moment for each of the 7 studios that offer FLF contracts for each instrument.

In all cases we hold the portfolio of titles fixed when contract types change[58]. Of course in reality the store may change both its portfolio of titles offered by this studio and the set of titles taken from other studios when it changes contract types: these are the market-expansion and leverage effects discussed above. However, we do not need to model the store's portfolio choices here in order to consistently estimate the inventory holding cost. The simpler inequalities that hold title portfolios fixed are also valid and are sufficient for our purposes. We model portfolio choices in the counterfactual analyses considered below.

[54] Some of this behavior may be explained by stores refusing to stock R rated titles, since FLF contracts can permit exceptions for objectionable content.There are also some clauses stating a maximum limit for the number of titles to be taken under FLF contracts per month. However, trial periods are a more likely explanation for these observations.

[55] We could alternatively define the start date as the first month of FLF after the end of the trial period, thereby excluding the initial period of learning about the new contract type, but this would reduce our number of FLF observations.

[56] If RS was not offered for the title, we switch it to a LP contract. Most stores take a FLF contract for at least 12 months. There are a few cases where the store drops a FLF contract after less than 12 months, either entering a test or trial program or exiting the contract early because of credit problems. We assume that the store expected the FLF contract to generate positive profits compared to its alternatives; if in reality store profits fell we assume this was due to unexpectedly high inventory costs (contained in our ε_{ms}). We assume that a move from FLF to RS contracts would not push the store's inventory above its maximum cutoff since the price per tape would increase and the store's profits per rental would fall. If a store ever goes on a FLF contract, then off, then on again, we reverse all of its FLF activity with the relevant studio.

[57] If this takes the store's inventory above the cutoff level we allow the store to reduce the number of tapes per title down to the minimum quantity restriction set by the studio, starting with box office category C titles (earliest-release date first, using descending order of inventory as a tie-breaker if necessary) and continuing until total inventory falls below the cutoff.

[58] The only portfolio change occurs when a store takes a new FLF contract: it then takes all titles offered by the studio in the following year, which may imply taking on new titles. In that case we define unobserved quality of the new titles to be the minimum ξ_{jmt} in the store-box-genre-month group. [As a robustness test we use the 25th percentile.] If there are no other titles in this group we use the average in the store-box-month.

In addition, in reality the store is making a dynamic choice. Because each title is active for five months, the optimal contract choice for a title released in month 1 should take into account both the effect of the choice on consumer demand for titles released in months 2-5 and also its effect on the store's choice of contracts for later titles. Our methodology accounts for the former but not the latter effect. Again this will not bias our estimates. The alternative contracts that we consider are feasible for the store; they may not be the optimal alternatives, since only a single contract choice is changed, but all we need is for the inequalities to hold in expectation and this requirement is still satisfied. Our assumption of single-firm markets helps us here: if more than one store existed in each market, each store would need to consider the future reactions of other agents when it changed its contract choice and the inequalities would need to take these predictions into account. Finally, it is worth noting that alternative methodologies such as the multinomial logit model would produce biased estimates when applied to this dynamic situation. This is one reason why we choose to use the inequalities methodology in preference to such alternatives.

The instruments $z_{ms'}$ (defined at the store-studio level) are required to be uncorrelated with $\varepsilon_{ms'}$, the unobservable in the profit equation, and correlated with the capacity chosen by the store. The unobservable includes variation in inventory holding costs and other store costs that is not observed by the store but that will affect its total costs when it alters its contracts with studio 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 studio in the relevant year and the percent of these titles that are from Box Office groups 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. The most troubling of these is a difference in the restrictions placed by studios on stores' sales of used tapes. 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 has to purchase a larger number of tapes, implying fewer rentals per tape, and is required to pay a certain proportion of the revenues from used tape sales back to the studio. This proportion may vary between RS and FLF contracts. This variation in requirements will bias our results if it affects stores' contract choices. However, 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 B and C titles In addition, the unobservable cannot include either econometrician measurement error or retailer prediction error in variables that affect demand or supply, or any effect of the heterogeneous outside option on the demand side that is not absorbed by the market fixed effects included there, because the nonlinearity introduced by the supply side quantity constraints implies that these unobservables will not be zeroed out when we take expectations conditional on the instruments. This is the reason for our assumption, noted above, that stores have perfect foresight regarding the demand side of the model.

Are there other possible endogeneity problems? We do not know of any other payments from store to studio or other store costs that are known to the store but that we do not observe in the data. It is possible that studios offer unobserved perks attached to choosing FLF contracts; however, these are likely to be minor items such as free posters for the store windows and are therefore unlikely to influence contractual choices. Stores may have credit constraints with a particular supplier that affect their contract choices: Rentrak sells tapes under RS and FLF contracts but those under

LP contracts are supplied by wholesalers, so if a store has credit problems with Rentrak it may be forced to choose a LP contract for all of its titles. The timing of stores' receipt of tapes could be important - for example, the choice of a FLF contract would be more attractive if the store received tapes sooner here than under LP/RS agreements. However, in reality all titles are shipped to stores on the release date, whatever the contract, and contract choices are made several months before this date. Similarly, we could be concerned about the endogeneity of $\bar{c}_{jm}$ and $\underline{c}_{jm}$, both of which are set by the studio on a title-by-title basis. Any unobservable that affects the store's choice of contract for title j may also affect the studio's choice of quantity restrictions. However, many studios choose these quantity restrictions using a formula based on the title's box office sales and the size of the store.

7.3 Results

We report here our preliminary results for the inequalities analysis. The moments generated by stores' choices between LP and RS contracts imply an estimated average cost of holding a tape, across all stores, of 72 cents. The 95% confidence interval for this estimate, calculated using the conservative confidence interval methodology in Pakes, Porter, Ho and Ishii (2006), is $[0.70, 0.87]$. This cost is approximately 26% of the average retail price per rental and a much smaller proportion of the cost to the store of purchasing the tape. As noted above, we plan to estimate a reduced form function for this cost, including store and studio characteristics in the specification. We will also repeat the estimation including inequalities generated from FLF contracts.

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. We would ideally separately identify the efficiency, market coverage and leverage effects. Our informational assumption for estimating the moment inequalities is that each store perfectly predicts all inputs to demand, including the titles to be released in future months, the store's choice of which titles to take and on which contract types, and consumer demand for these titles. We require our counterfactual experiment to be consistent with this assumption.

We will use our estimates to predict the effect of postponing the implementation of FLF contracts for a subset of the studios in our data. Two of the seven FLF studios implemented FLF early, in month 14 of our panel. The others followed some time later: in or after month 30. We will leave the two early implementers unchanged, assuming that any store learning about FLF contracts will therefore be completed before our counterfactual begins. We will then remove the three earliest months of FLF for two studios: those that began FLF in months 30 and 31 respectively. (We will hold trial periods fixed and remove just the first three months in which no stores undertook trials.) For each store we will solve backwards from the last to the first month from which FLF has been removed, considering the choice of which titles to take for all studios. Our assumption that each store operates in a separate market implies a unique equilibrium. We will obtain a set of sequential choices over titles' contract types that are optimal for the store assuming perfect foresight, given the observed choices in the data both before and after the three-month counterfactual period.

We propose to hold fixed stores' choices of contract types both before and after our counterfactual. We will assume that the removal of FLF was a surprise, so choices made before the removal may not be optimal after the change. The assumption regarding the period after the counterfactual is more difficult to justify. In reality, 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 propose two alternative methods to address this issue. First we will assume that stores make their choices in the belief that FLF will be available from the end of the trial periods onwards. The FLF studios then unexpectedly delay the beginning of the program by three months but hold stores to their commitments to purchases after FLF is reinstated. Other studios also hold stores to these commitments but allow them to re-optimize during the counterfactual period. This latter assumption is probably unrealistic; its effect on our results depends on the extent to which choices in one month are influenced by expected contract types in later periods. Our second counterfactual will address this by assuming that the non-FLF studios 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 counterfactual analyses which indicate that the leverage effect is probably small).

Our results will inform us, first, about the number of titles taken by each store from each studio after the removal of FLF contracts. Comparing these estimates to the observed data will provide information on the market coverage and leverage effects. We will also use the estimates to predict the change in consumer surplus, store profits and studio revenues per month resulting from the removal of FLF contracts. Finally we will note what can be concluded about the reasons why studios choose to offer these contract types.

9 Discussion and Conclusion

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.

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* Vol. 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, 13, 185-213.

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

Chu, C., Leslie, P. and Sorenson, A., (2007), "Nearly Optimal Pricing for Multiproduct Firms," *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* Vol. 122, pp. 1307-50.

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

Conlon, C. and Mortimer, J., (2007), "Demand Estimation Under Incomplete Product Availability", *working paper.*

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

Pakes, Porter, Ho and Ishii, (2006), "Moment Inequalities and Their Application", working paper.

Posner, (1975), "The Social Costs of Monopoly and Regulation," *Journal of Political Economy* Vol. 83, No. 4, pp 807-27.

Salinger, (1995), "A Graphical Analysis of Bundling," *Journal of Business* Vol. 68, No. 1, pp 85-98.

Schmalensee, (1982a), "Antitrust and the New Industrial Economics," *American Economic Review, Papers and Proceedings* Vol. 72, No. 2, pp 24-28.

Schmalensee, (1982b), "Commodity Bundling by Single-Product Monopolies," *Journal of Law and Economics* Vol. 25, No. 1, pp 67-71.

Whinston, (1990), "Tying, Foreclosure and Exclusion," *The American Economics Review* Vol. 80, No. 4, pp 837-59.

Figure 1: Elasticities Implied by Demand Estimates

Average Price Elasticities

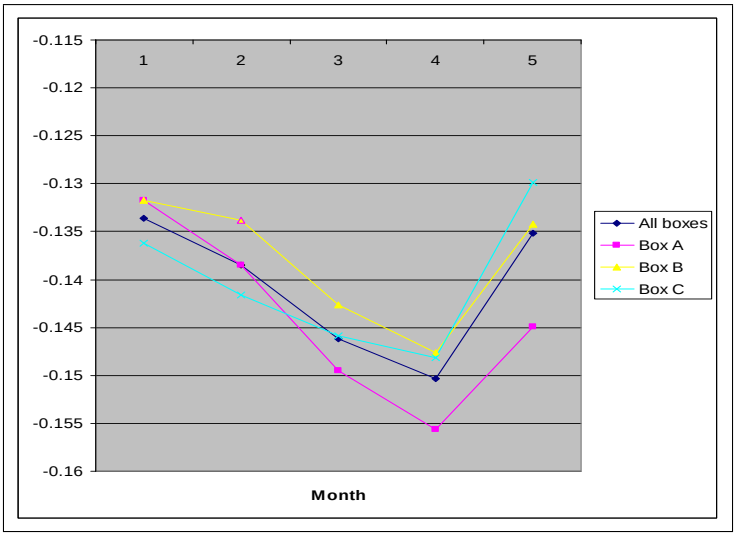

Average Inventory Elasticities

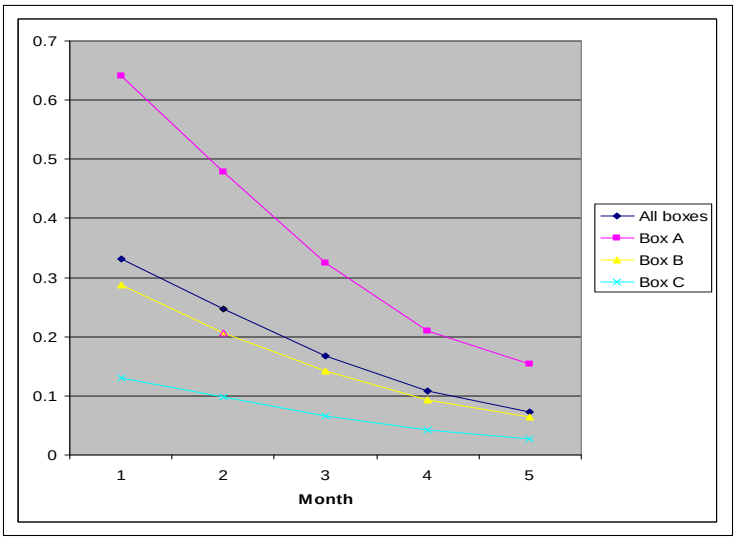

Notes: Elasticities implied by the demand estimates. Corresponding data are given in Table 6. Elasticities with respect to price and inventory are calculated for every store-title-month observation and then averages are taken within each zipcode category-month and then across zipcode categories.

Table 1: Summary Statistics

Contract	Linear Pricing	Revenue Sharing	Full-Line Forcing	Sell-Through Priced
Avg Terms				
Upfront Fee	66.82	8.48	3.60	15.17
	(5.59)	(1.07)	(1.24)	(1.64)
Retailer's Share of Revenue	100%	45.95%	59.01%	100%
	(-)	(2.98%)	(2.00%)	(-)
Minimum quantity	-	9.64	10.68	-
	(-)	(11.47)	(10.58)	(-)
Maximum quantity	-	23.37	22.28	-
	(-)	(22.59)	(21.44)	(-)
Avg No. of Rentals				
Month 1:	49.58	67.46	52.13	91.01
	(78.47)	(100.7)	(88.39)	(125.5)
Month 2:	64.87	70.38	60.01	83.97
	(90.21)	(100.3)	(80.54)	(100.2)
Month 3:	38.87	35.40	32.39	39.22
	(51.57)	(49.06)	(44.18)	(46.38)
Month 4:	25.13	22.26	20.48	21.94
	(31.77)	(29.73)	(26.73)	(25.16)
Month 5+:	69.54	60.06	57.29	78.73
	(102.8)	(88.53)	(81.13)	(136.0)
Avg Rental Price				
Month 1:	2.67	2.69	2.71	2.71
	(0.61)	(0.52)	(0.63)	(0.58)
Month 2:	2.85	2.80	2.90	2.89
	(0.61)	(0.56)	(0.60)	(0.63)
Month 3:	2.85	2.80	2.89	2.96
	(0.66)	(0.60)	(0.67)	(0.75)
Month 4:	2.85	2.80	2.88	2.97
	(0.70)	(0.64)	(0.70)	(0.86)
Month 5+:	2.82	2.72	2.91	2.98
	(0.74)	(0.70)	(0.76)	(0.90)
Avg Rentals per Tape				
Month 1:	5.63	4.27	4.13	5.13
	(4.42)	(2.89)	(3.13)	(4.73)
Month 2:	7.73	4.78	5.38	5.20
	(4.92)	(3.22)	(3.71)	(3.74)
Month 3:	5.06	2.53	3.38	2.60
	(3.78)	(1.89)	(2.85)	(2.14)
Month 4:	3.56	1.66	2.45	1.61
	(2.94)	(1.33)	(2.41)	(1.66)
Month 5+:	13.52	5.04	7.28	6.82
	(14.32)	(5.07)	(9.00)	(8.84)
Avg Inventory	8.81	14.60	12.53	18.20
	(13.92)	(17.64)	(17.33)	(21.73)

Notes: Average contract terms, rentals, prices, inventories and store sizes for titles taken under each contract type. Averages are across store-title pairs. Standard deviations in parentheses. Data source: Rentrak Corporation, January 1, 1998 to June 30, 2002.

Table 2: Summary Statistics (Cont.)

Contract	Linear Pricing	Revenue Sharing	Full-Line Forcing	Sell-Through Priced
Total No. of Titles Offered by Studios				
Year 1:	219	150	0	27
A Titles:	30	24	0	15
B Titles:	36	25	0	6
C Titles:	153	101	0	6
Year 2:	195	143	10	24
A Titles:	32	30	1	14
B Titles:	40	36	2	6
C Titles:	123	77	7	4
Year 3:	221	173	18	21
A Titles:	41	39	4	15
B Titles:	43	39	3	1
C Titles:	137	95	11	5
Year 4:	195	122	39	26
A Titles:	33	20	9	16
B Titles:	49	23	5	3
C Titles:	113	79	25	7

Notes: Average number of titles offered by studios 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 terms. No Sell-through Priced titles are offered under other contracts. The primary source of data summarized in this table is Rentrak Corporation. Data were gathered from January 1, 1998 to June 30, 2002.

Table 3: Summary Statistics (Cont.)

Contract	Linear Pricing	Revenue Sharing	Full-Line Forcing	Sell-Through Priced
Number of Stores				
	7107	6687	4896	6926
Avg No. of Titles Taken by Stores				
Year 1:	95.86	23.02	-	19.31
	(45.20)	(24.76)	-	(7.81)
A Titles:	19.23	4.88	-	11.58
	(7.70)	(3.58)	-	(4.45)
B Titles:	22.92	4.92	-	4.06
	(10.08)	(5.42)	-	(1.89)
C Titles:	53.71	13.21	-	3.68
	(30.18)	(16.95)	-	(2.00)
Year 2:	68.36	20.11	3.64	14.90
	(44.62)	(21.68)	(3.23)	(7.95)
A Titles:	14.25	7.01	0.37	8.82
	(8.86)	(6.13)	(0.48)	(4.41)
B Titles:	19.67	6.76	1.06	3.85
	(11.90)	(7.61)	(0.96)	(2.42)
C Titles:	34.43	6.34	2.21	2.24
	(25.95)	(9.28)	(2.22)	(1.61)
Year 3:	95.63	18.35	6.04	14.23
	(55.63)	(23.37)	(3.92)	(6.93)
A Titles:	23.05	7.30	1.53	11.09
	(12.38)	(7.31)	(1.01)	(5.31)
B Titles:	25.43	5.23	0.86	0.74
	(13.69)	(7.40)	(0.67)	(0.44)
C Titles:	47.15	5.82	3.65	2.40
	(31.88)	(9.96)	(2.69)	(1.49)
Year 4:	85.38	13.79	10.01	17.22
	(40.80)	(17.41)	(7.35)	(7.51
A Titles:	19.60	5.56	3.27	12.11
	(9.10)	(5.32)	(2.10)	(5.02)
B Titles:	32.07	3.54	0.97	2.06
	(13.77)	(5.27)	(0.99)	(1.30)
C Titles:	33.71	4.69	5.77	3.05
	(20.41)	(7.73)	(4.69)	(1.64)

Notes: Average number of titles of each contract type taken by all active stores in each year, where stores not participating in a contractual form in a given year are listed with zero titles. Standard deviations in parentheses. The primary source of data summarized in this table is Rentrak Corporation. Data were gathered from January 1, 1998 to June 30, 2002.

Table 4: Summary Statistics (Cont.)

Contract		Linear Pricing	Revenue Sharing	Full-Line Forcing	Sell-Through Priced
Ave store tier					
Quintile 1		2.52	5.24	2.91	4.75
Quintile 2		3.07	4.94	4.18	4.64
Quintile 3		3.46	3.16	4.47	4.03
Quintile 4		4.83	3.01	4.01	2.91
Quintile 5		5.11	2.70	4.54	2.31
Quintile	% of quintile				
1	Tier 1-3	1.35	0.22	1.19	0.62
	Tier 7-10	0.20	1.13	0.24	1.50
2	Tier 1-3	1.13	0.50	0.63	0.48
	Tier 7-10	0.33	1.39	0.60	0.99
3	Tier 1-3	0.99	1.10	0.57	0.68
	Tier 7-10	0.50	0.38	0.93	0.43
4	Tier 1-3	0.38	1.16	0.79	1.17
	Tier 7-10	0.88	0.33	0.82	0.38
5	Tier 1-3	0.43	1.30	0.62	1.41
	Tier 7-10	1.56	0.42	1.06	0.09

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 size in each quintile. Tiers are ranked from 1 to 10 where 10 is largest. Panel 2 reports the percent of stores in each quintile of percent of titles adopted under a particular contract type that are in store sizes 1-3 and 7-10 respectively. These numbers are normalized by the percent of all stores that are in the relevant set of tiers. Numbers over 1 indicate that this type of store is over-represented in the relevant quintile for this contract type. The primary source of data summarized in this table is Rentrak Corporation. Data were gathered from January 1, 1998 to June 30, 2002.

Table 5: Demand Results

	OLS	IV 1	IV 2
	Coefft (S.E.)	Coefft (S.E.)	Coefft (S.E.)
Price	-.027 (.002)	-.026 (.003)	-.024 (.003)
Month 2	.131 (.022)	.154 (.024)	.168 (.025)
Month 3	-.137 (.022)	-.191 (.025)	-.172 (.025)
Month 4	-.399 (.023)	-.505 (.025)	-.478 (.026)
Month 5+	.192 (.024)	.277 (.027)	.313 (.027)
Inventory	.0187 (.0003)	.0210 (.0004)	.0192 (.0004)
Inv*Month 2	-.003 (.0004)	-.004 (.0005)	-.005 (.0005)
Inv*Month 3	-.008 (.0004)	-.009 (.0005)	-.011 (.0005)
Inv*Month 4	-.012 (.0004)	-.013 (.0005)	-.015 (.0005)
Inv*Month 5	-.011 (.0004)	-.013 (.0004)	-.016 (.0005)
σ	.631 (.002)	.497 (.003)	.499 (.003)
N	405831	405831	405831
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 group, 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 groups:					
Price elasticity	-0.134	-0.138	-0.146	-0.150	-0.135
Inventory elasticity	0.331	0.246	0.168	0.108	0.072
Box Office Group A:					
Price elasticity	-0.132	-0.138	-0.150	-0.156	-0.145
Inventory elasticity	0.640	0.479	0.325	0.210	0.153
Box Office Group B:					
Price elasticity	-0.132	-0.134	-0.143	-0.148	-0.134
Inventory elasticity	0.288	0.207	0.142	0.093	0.065
Box Office Group C:					
Price elasticity	-0.136	-0.142	-0.146	-0.148	-0.130
Inventory elasticity	0.131	0.098	0.066	0.042	0.028

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 first within each zipcode category-month and then across zipcode categories.

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)	-0.011 (0.004)	-0.014 (0.004)
Quarter 3	-0.118 (0.004)	-0.120 (0.004)	-0.122 (0.004)
Quarter 4	-0.029 (0.004)	-0.041 (0.004)	-0.042 (0.004)
Box Office:			
B	-0.745 (0.042)	-0.685 (0.044)	-0.701 (0.046)
C	-1.360 (0.039)	-1.277 (0.041)	-1.291 (0.042)
Genre:			
Child/Family	-0.772 (0.051)	-0.589 (0.052)	-0.605 (0.052)
Comedy	0.187 (0.046)	0.277 (0.047)	0.270 (0.047)
Drama	-0.133 (0.023)	-0.068 (0.073)	-0.077 (0.023)
Horror/Suspense	-0.021 (0.031)	-0.023 (0.033)	-0.027 (0.034)
Romance	-0.814 (0.045)	-0.657 (0.046)	-0.669 (0.047)
Science Fiction	-0.722 (0.053)	-0.551 (0.056)	-0.545 (0.055)
Rating:			
PG	-0.020 (0.032)	0.027 (0.032)	0.033 (0.033)
PG13	0.067 (0.053)	0.183 (0.054)	0.190 (0.054)
R, NC17, NR	0.011 (0.059)	0.130 (0.060)	0.131 (0.061)
Market characteristics:			
Median age	0.007 (0.001)	0.008 (0.001)	0.006 (0.001)
Median income	-0.011 (0.0002)	-0.011 (0.0002)	-0.012 (0.0002)
Number of households	-7.1e-5 (3e-7)	-7.2e-5 (4e-7)	–7.1e-5 (3e-7)
Percent white	-0.029 (0.0004)	-0.029 (0.0004)	-0.030 (0.0005)
Percent black	-0.024 (0.0004)	-0.023 (0.0004)	-0.024 (0.0004)
Percent female	0.016 (0.002)	0.013 (0.002)	0.019 (0.002)
Percent single mother with kids	0.006 (0.001)	0.006 (0.001)	0.007 (0.001)
Percent single father with kids	-0.258 (0.004)	-0.254 (0.004)	-0.255 (0.004)
Percent single male	-0.374 (0.008)	-0.375 (0.008)	-0.372 (0.008)
Percent single female	-0.034 (0.004)	-0.037 (0.004)	-0.035 (0.004)
Percent married with kids	0.020 (0.001)	0.019 (0.001)	0.022 (0.001)
Percent with Bachelor's	-0.012 (0.0003)	-0.012 (0.0003)	-0.011 (0.0003)
Blockbuster in market	0.670 (0.003)	0.670 (0.004)	0.674 (0.004)
Percent rural	0.005 (0.0002)	0.0004 (0.0002)	0.0006 (0.0002)
Percent suburban	-0.001 (0.0001)	-0.001 (0.0001)	-0.001 (0.0001)
N	405831	405831	405831
R^2	0.43	0.40	0.39

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.