Digital Distribution and the Prohibition of Resale Markets for Information Goods

Benjamin Reed Shiller, Economics Department, Brandeis University

Working Paper Series

2013 | 58

Digital Distribution and the Prohibition of Resale Markets for Information Goods

Benjamin Reed Shiller*
Brandeis University

Dec, 2012

Abstract

An existing theoretical literature finds that frictionless resale markets cannot reduce profits of monopolist producers of perfectly durable goods. This paper starts by presenting logical arguments suggesting this finding does not hold for goods consumers tire of with use, implying the impact of resale is an empirical question. The empirical impact is then estimated in the market for video games, one of many markets in which producers may soon legally prevent resale by distributing their products digitally as downloads or streamed rentals. Estimation proceeds in two steps. First, demand parameters are estimated using a dynamic discrete choice model in a market with allowed resale, using data on new sales and used trade-ins. Then, using these parameter estimates, prices, profits, and consumer welfare are simulated under counterfactual environments. When resale is allowed, firms are unable to prevent their goods from selling for low prices in later periods. The ability to do so by restricting resale outright yields significant profit increases. Renting, however, does not raise profits as much due to a revenue extraction problem. (JEL M30, L00, K19)

1 Introduction

The standard theoretical model of perfectly durable goods finds that frictionless resale does not lower monopolist profits, implying producers would welcome well-functioning resale markets for their products. The reasoning behind this finding is that forward-looking consumers incorporate the expected resale price into initial willingness to pay, and the increase in revenue from higher initial prices offsets revenue lost from fewer sales. Since firms may earn the same revenues with fewer goods produced, profits may be higher when resale is allowed.

*I owe much gratitude to Joel Waldfogel, Katja Seim, and Alon Eizenberg for numerous discussions, suggestions, and advice. I would also like to thank Ulrich Doraszelski and Lorin Hitt for their very helpful input, Tim Derdenger for advice at different stages, and Shane Greenstein, Adam Isen, Phil Leslie, David Muir, Felix Oberholzer-Gee, Andrew Paciorek, John Riley, David Rothschild, Kent Smetters, Rahul Telang, and Jeremy Tobacman for their suggestions. I want to thank Hamilton Chu, David Edery, and Marc Mondhaschen for imparting to me some of their wisdom on the video game industry. I am grateful for support received from the NBER as a visiting fellow. I appreciated useful comments and suggestions from conference participants at Brandeis University, Carnegie Mellon University, Georgia Institute of Technology, Microsoft Research, NBER Summer Institute, UCLA, and the University of Pennsyvlania.

If one modifies the standard model of resale by allowing consumers to tire of goods, however, this finding no longer holds, implying the impact of resale on profits is ambiguous. The intuition is as follows. Owners, having grown tired with use, resell their used products in later periods. These used goods fulfill some residual demand, clearing the market at a low price. Anticipating the low price and extant high quality to someone that has not used it, forward-looking consumers are not willing to pay as much as they otherwise would up front.

The assumption that consumers tire of products with use is relevant for many types of information goods. It clearly applies to books, movies, video games, and music, which together yield about $90 billion in sales in the U.S. alone.[1] It also applies to some non-entertainment products like learning software (e.g. languages) and may even apply to more traditional durables if consumers are excited by initial use.

For many such goods, resale markets exist not because they weakly raise producer profits - nay, practitioners strongly believe otherwise - but because there has been no legal way to prevent them. The U.S. first sale doctrine (17 U.S.C. section 109) guarantees consumers' right to resell the original purchased copy of a durable good, even if copyrighted, so long as no copies were made.[2]

However, information good producers may soon be able to prevent resale legally. Because transferring a digital file involves making a copy from the hard drive, and the first sale doctrine only applies to the original copy, the doctrine does not apply to digital downloads.[3,4,5] Firms can also prevent resale indirectly by streaming rentals from places like Netflix and Spotify. Moreover, with the advent of high speed internet connections and adoption of platforms capable of preventing illegal file sharing, e.g. Kindle, exclusive digital distribution is becoming more feasible. A case is set to be decided by the Supreme Court to determine whether firms can prevent resale of physical goods as well by manufacturing overseas.[6]

To investigate the empirical impact of shutting resale markets in the video game market, I employ a two-step approach. First demand parameters are estimated in a market where resale exists, using a dynamic discrete choice structural model of the consumer's purchase and resale decisions and a dataset containing new video game purchases and

[1]See Rich (2009), Hurt (2009), Digital Entertainment Group's "Digital Entertainment Group Year-End 2008 Home Entertainment Sales Figures," and the International Federation of the Phonographic Industry's "Recorded Music Sales 2008" report.

[2]The first-sale doctrine dates back to an 1854 Supreme Court case, Stevens v. Royal Gladding, which ruled that a cartographer's right to sole distribution ended at first sale. It was subsequently codified in 1909, and updated in 1976. It has been unclear whether the first-sale doctrine applies to licensed goods. A District Court judge decided in Vernor v. Autodesk in 2008 that permanent licenses constituted sales, and as a result were covered by the first-sale doctrine. In September 2010, the Ninth Circuit Court of Appeals reversed this ruling, determining that firms can legally prohibit resale in their licensing agreement even if they never intended for the good to be returned to them. In 2011, the ninth circuit took the opposite view in Universal Group v. Augusto. Capital records v. ReDigi, currently ongoing, considers whether licensed digital songs can be resold.

[3]To resell a downloaded good, one would need to sell the hard drive that contains it. In many cases, this would mean selling the entire device, along with all other information goods downloaded to it. See Graham (2002), Hinkes (2007), Long (2007), and Seringhaus (2009).

[4]Firms can enforce prohibited resale with access control software, and following the Digital Millennium Copyright Act (1998) can prosecute creators of software designed to circumvent copyright protection software.

[5]In the Balance Act of 2003, Congress considered instituting a digital first-sale doctrine, allowing resale of digital goods via the "forward and delete" resale method. However, the bill did not pass.

[6]Kirtsaeng v. John Wiley and Sons, Inc.

used game resales. Second, using the parameter estimates, profits are simulated in counterfactual environments.

The parameter estimates show that consumers tire of video games very quickly with use. High valuation consumers reduce their value, relative to the outside good, from $80 on average per month initially to just a couple of dollars per month by the sixth month. As a result, resale markets put downward pressure on price. In counterfactual simulations, I find that optimal prices fall at a much slower rate when resale markets are shut down, providing much less incentive to delay purchase. Selling non-resellable downloads, rather than resellable versions, is estimated to raise producer profits substantially, by 109%. However, the magnitude of this finding is sensitive to the assumed market size. Preventing resale indirectly by exclusively renting to consumers is not nearly as attractive a strategy - it only raises profits by 5% percent.

This paper contributes to a sparse empirical literature on resale and monopolist profits. Chen et al. (2011) estimate the impact of resale in the market for cars, finding that resale would raise a monopolist producer's profits by 15%.[7] However, since they investigated cars, they did not allow for consumers to tire with use. A contemporaneous paper by Ishihara and Ching (2012) allows for such tiring. They focus on the Japanese video game market, which has some very different characteristics and norms, such as non-declining prices for new games.[8] To appropriately model Japan, they assume novelty effects which reduce usage value as the game ages for consumers yet to play the game, but do not allow for heterogeneous usage values. Despite these differences, they find results that are qualitatively similar to the findings in this paper, although their results are less strong. Focusing on airline tickets, Lazarev (2012) also finds resale would lower profits, by 29%.[9] To my knowledge, no papers investigate the empirical impact of preventing resale via renting.

In the next section, I discuss the assumption of losing interest with use, and contrast the logic of this assumption with prior theory. A basic industry background is also provided. Section 3 describes the data. Then, in sections 4 and 5, the model of consumer demand and the estimation strategy are detailed. Section 6 presents the results.

2 Background

2.1 Logic and Prior Theory

This subsection presents logical arguments illustrating why monopolist revenues from selling a durable good may be higher when resale is prohibited compared to when resale is frictionless, when one assumes consumers tire of products with use. The same logic suggests rental profits may be less than non-resellable good profits. To show this, I first argue the second point, that producer revenues from selling a non-resellable good are higher than revenues from renting directly to consumers. Next it is shown that, by extension, this implies producer revenues under frictionless resale are also lower than when resale is prohibited. This logic is then compared to the prior theoretical literature.

When forward-looking consumers tire of the good, renting may be unable to yield monopoly revenues as high as selling non-resellable goods. To illustrate this point, suppose a single individual. When the firm sells a non-resellable good up front it can success-

[7] For oligopolist producers, Chen et al., (2011) find resale lowers profits.

[8] By contast, in the U.S. prices fall by roughly half in the first year following release.

[9] Leslie and Sorensen (2011) also investigate impact of ticket resale on consumer welfare.

fully charge the individual's full expected discounted utility, if it commits to maintaining high prices. For a rental pricing strategy to extract full surplus, the firm would need to charge the consumer a price each period equal to her full usage value for that period, implying a declining rental price. However, the firm could not successfully rent the product each period at such prices, for the following reason. If prices decline, the consumer would prefer to delay renting until prices were lower thereby yielding positive, rather than zero, consumer surplus. Hence, the separate per-period prices introduce an incentive compatibility constraint not present when selling a non-resellable good. The underlying reasoning extends beyond a single individual, suggesting monopoly revenue when renting may be lower than when selling a non-resellable good.

Revenues under frictionless resale may be lower than revenues from non-resellable goods for the same reason. When there is frictionless resale, a rental price is implied by the change in prices period to period. A declining implicit per-period rental price introduces an incentive compatibility constraint, limiting how much forward-looking consumers will pay for use of the product. Preventing resale allows the firm to charge a flat price to consumers for permanent ownership, removing this incentive compatibility constraint, and allowing them to extract a greater share of the surplus.

Other logic suggests that resale and renting may raise or alternatively lower profits. First, since goods change hands under resale and renting, the same stream of services might be provided with fewer produced goods. Hence, the firm may have equal or similar revenues but lower costs when selling a resellable good. Second, frictions in the resale market may substantially lessen competition from used goods and hence reduce the rate at which prices decline, loosening the incentive compatibility constraint. Third, resale can prevent profit-increasing intertemporal price discrimination (Lazarev, 2012).

The above arguments imply that selling a non-resellable good may or may not yield higher profits than renting or selling a resellable good. The relative profits of these three strategies is hence an empirical question.

By contrast, existing models of secondary markets for durables, which do not allow consumers to tire of goods, typically find that frictionless resale can raise, but not lower, monopolist producer profits (Hendel and Lizzeri, 1999; Rust, 1986).[10,11] This finding is common for both depreciating durables and One-Hoss Shay durable goods, a category that includes information products.[12,13,14]

At first glance, it might appear that quality depreciation functions similarly to growing tired of a product. This is not so. When consumers tire of non-depreciating goods, they can obtain the full quality product for a lower price by waiting. This implies consumers have a strong incentive compatibility constraint limiting how much they will pay. For goods that depreciate, however, waiting to buy a used good entails obtaining a lower quality product, generally offsetting the gain from obtaining the product at a lower price.

[10]Bulow (1982), Hendel and Lizerri (1999), and Rust (1986) show that imperfect durability on its own may lower profits, because owners have the option of keeping a good they purchased earlier rather than returning to the market to buy a new product.

[11]Resale can lower producer profits if resale transaction costs exceed primary market transaction costs.

[12]One-Hoss Shay durables have the property of random binary depreciation - they work fully until they don't work at all. A common example is light bulbs.

[13]Papers that find resale can lower firm profits (Anderson and Ginsburgh, 1994; Miller, 1974) generally still find that resale markets do not lower monopolist producer profits when goods are perfectly durable.

[14]Preventing resale may raise profits when demand fluctuates substantially, and goods change hands each period (e.g. textbooks). See the hide/steers analogy in Benjamin and Kormendi (1974).

2.2 Industry Background

When purchasing a game, consumers can generally choose between a new and used copy, and between online and offline sellers. A used copy sells to consumers for close to the same price as a new copy. The small price difference (10%) is usually assumed to be due to risks involved in buying a used copy which may not function, and negative feelings associated with buying a used copy experienced at time of purchase. However, as long as the used game is not sufficiently damaged (such as having a bad scratch), it provides exactly the same service as a new copy. It also is typically assumed that the larger price difference between copies bought in brick and mortar stores and via auction, both for new and used games, reflects additional costs borne by the buyer in an auction. Such costs include shipping fees, risk of being scammed, the cost of postponing gratification while the auction concludes and the game ships, etc. It seems reasonable to assume that the true total cost of buying a copy of a game does not depend on the condition, new or used, or whether bought online. Hence, consumers view these choices as roughly equivalent in value.

Consumers, when they have tired of a game, have the option to resell their copy either directly back to another consumer via auction, or to a retailer. It seems logical that since these different mediums are competing for traded-in games, and consumers can freely choose where to resell their games, that the price they receive for trade-ins does not heavily depend on where they choose to resell it.

The used game market was generally considered concentrated, but the market for new games was not. One company, GameStop, was and generally still is regarded as the dominant buyer and seller of used games. In private correspondence, the former director of used games at GameStop's main brick and mortar competitor at the time, GameCrazy, assured that he had estimated GameStop's used game market share to be 80% in 2010.[15] In addition to GameCrazy, GameStop also competed in the resale market with online auction websites like eBay and niche online retailers. Sales of new games, however, were less concentrated. Based on calculations explained later, GameStop was estimated to account for approximately 25% of new sales, and competitors such as Best Buy, Target, and Walmart also had substantial new game sales market share. While these major retail chains contemplated entering the used game market in earnest, none did until 2009, which is after the period investigated in this paper.

See Lee's (2012a) review paper for a more thorough description of video game platforms and the market for new copies of video games, as well as a discussion of related literature.

3 Data

The data used in this paper were constructed from two datasets. The first dataset, from the NPD group, provides information on monthly sales and average prices of new copies of XBOX 360 video games by game in the U.S. from November 2005 to December 2008.[16] The second dataset contains used game auctions from a popular online marketplace over that same timespan. These latter data serve as a proxy for the monthly quantity of used

[15] GameCrazy shared retail space with Hollywood Video. When its parent company, Movie Gallery, declared bankruptcy in 2010, Hollywood Video stores, and hence Game Crazy stores, were closed.

[16] NPD observes over 80% of point of sales transactions of video games and scales them up to the market.

games *traded-in* by consumers, i.e. resold by consumers to retail stores. These data also contain prices consumers receive for reselling.[17] Because these auctions will proxy for trade-ins, I will subsequently refer to these data as the trade-ins data. Time-invariant variables in the combined dataset include game characteristics such as composite critic review scores and genres from the NPD group and "replay value" scores from Game Informer Magazine's reviews.

The trade-ins data, which comprise only a share of the U.S. market, must be scaled-up to the entire U.S. market to be consistent with the data on sales of new copies of games. The most obvious method for scaling-up the trade-ins to reach an estimate of total market trade-ins, by multiplying them by the inverse of the auction seller's used game market share, is not feasible for two reasons. First, the used game market share is not available. Sales are not revealed by most firms nor collected by consulting companies like NPD. Second, market share information would be insufficient, since the data omit an unknown number of auctions.[18] However, the scale-up factor can be computed another way.

Suppose one had information on the ratio of cumulative sales of used games by retailers to cumulative sales of new copies within s time periods following each game's release. This ratio, labeled $Ratio^M$, is shown below:

$$Ratio^M = \frac{\sum_j Q^{UsedSales}_{j,t-r\leq s}}{\sum_j Q^{NewSales}_{j,t-r\leq s}} \tag{1}$$

where $Q^{UsedSales}_{j,t-r\leq s}$ is the number used copies of game j sold from retailers back to consumers within s periods of the game's release date r, and $Q^{NewSales}_{j,t-r\leq s}$ is the corresponding number of new copy sales. If this ratio were known, one could estimate the total cumulative used sales of these games by all retailers to consumers within s periods of release by multiplying $Ratio^M$ by the corresponding total cumulative new copy sales. After one modification, explained later, the trade-ins data could then be scaled up to match, thus proxying for the entire market. The corresponding scale-up factor can then be used to estimate trade-ins and used sales by retailers at other lengths of time following a game's release, i.e. other than s.

Note the distinction between used sales by retailers to consumers, and trade-ins by consumers to retailers. The auction data are a good indicator for the decision to resell, and hence a good proxy for used game trade-ins by consumers back to the retail store, how most games are resold. They, however, are not a good proxy for sales of used games by a retail store back to consumers, because they do not account for the length of time used games are held in inventory. Hence, the cumulative number of trade-ins should exceed cumulative used sales out of inventory if there is any meaningful inventory lag. Since $Ratio^M$ applies to used sales by retailers back to consumers, and does not include unsold inventory of traded-in games, the trade-ins need to be converted to expected used sales by retailers to consumers before applying the scale-up procedure.

One can estimate the cumulative number of used copy sales of game j from inventory

[17] Consumers had several outlets for reselling games, suggesting a competitive market price received for resold (traded-in) games

[18] The data omit an unknown number of auction sales with missing platform identifiers. Additionally, the dataset does not classify the condition (new/used) for all auction sales. 13.5% of observations in the data were missing condition (new/used) values. The method from Gentzkow and Shapiro (2010) for identifying media slant keywords was used to reclassify 22% of these missing condition values to new or used.

that would have occurred at a retail store if the number and timing of their trade-ins matched the data. Label this as $Q_{j,t-r\leq s}^{UsedSalesSample}$. I assume the length of time, in days, from trade-in to subsequent resale by the retailer to another consumer follows the geometric distribution with probability p equal to 0.015748, the inverse of GameStop's average inventory time to sale.[19] The cumulative used sales by this retailer to consumers approximately equals:

$$Q_{j,t-r\leq s}^{UsedSalesSample} = \sum_{d=r}^{s+r} \sum_{t'=d+1}^{s+r} Q_{j,d}^{TradeInsSample} * p * (1-p)^{t'-d} \tag{2}$$

where $d-r$ equals the days since release, and $Q_{j,d}^{TradeInsSample}$ is the number of trade-ins of game j by consumers on the $(d-r)^{th}$ day following release. It is estimated that about a third of trade-ins over this time period are subsequently resold by retailers back to other consumers.

$Ratio^M$, necessary to compute the scale-up, is not directly available but can be computed. Note that if one multiplies $Ratio^M$ by the corresponding ratio for a single retailer and this latter ratio's inverse, one can rearrange to yield the following equation:

$$Ratio^M = \frac{\sum_j Q_{j,t-r\leq s}^{NewSales,GS} / \sum_j Q_{j,t-r\leq s}^{NewSales}}{\sum_j Q_{j,t-r\leq s}^{UsedSales,GS} / \sum_j Q_{j,t-r\leq s}^{UsedSales}} * \frac{\sum_j Q_{j,t-r\leq s}^{UsedSales,GS}}{\sum_j Q_{j,t-r\leq s}^{NewSales,GS}} \tag{3}$$

where $\frac{\sum_j Q_{j,t-r\leq s}^{UsedSales,GS}}{\sum_j Q_{j,t-r\leq s}^{NewSales,GS}}$, which includes the superscript GS, is the corresponding ratio for a single retailer. According to equation 3, $Ratio^M$ can be computed by multiplying the corresponding ratio for a single retailer by the ratio of that retailer's new copies market share to used copies market share.

The information necessary for equation 3 is available for the market leader, GameStop. GameStop's domestic new game revenues are estimated by multiplying their global new software sales revenue by the share of their stores that are in the U.S.[20] Dividing GameStop's U.S. revenues from new copy sales by total U.S. new game sales revenues yields GameStop's estimated share. The three year average over 2006-2008 is 24.80%.[21] GameStop's used game market share is not well documented. However, the former Director of Used Games at GameStop's biggest competitor estimated, based on internally gathered data, that GameStop's market share was 80% in early 2010. The last component needed, the ratio of cumulative used to new game sales at GameStop within some time of release can be derived from information revealed by the Vice President of Marketing. He stated that 4% of GameStop's total sales, new plus used, within the first two months following games' releases are of used copies. This implies the ratio of cumulative used to new sales over that time frame equals $4.1\overline{6}\%$.

The scale-up factor can be found from the above information. First note that, following equation 3, $Ratio^M = 24.80/80.00 * 0.041\overline{6} = 0.0129$. This implies that for each 100 copies of a new game that are sold in the first two months following release, on average 1.29 used copies will be sold by retailers back to consumers. Multiplying by cumulative

[19] According to GameStop's financial disclosures, the three year mean, 2006 to 2008, of their average annual inventory turnover is $5.\overline{6}$. This implies average time to sale of inventory of 63.5 days.

[20] The share of GameStop's stores in the U.S. closely tracks the share of their total revenues (also including hardware sales) that are in the U.S. The domestic store shares (revenues shares) are 0.78 (0.80) in 2006, 0.77 (0.77) in 2007, and 0.70 (0.73) in 2008.

[21] Market share estimates over the three years are very steady. They are: 24.85%, 24.85%, and 24.69%.

new sales in the entire market in the first two months following each game's release yields an estimate of about $647,000$ sales of used copies by retailers to consumers in the same time-frame. Scaling the data's implied used sales by retailers to consumers - calculated from trade-ins according to equation 2 - to match the $647,000$ expected used sales yields a scale-up factor equal to 42.83. The implied market share of the trade-ins data is 2.3%.

One more adjustment to the data is required before proceeding. The model requires information on the total number of copies of a game sold by retailers to consumers, including both new and used copies. The number of new games sold is directly available from the NPD dataset. The sales of used copies by retailers to consumers is estimated from the trade-ins data, assuming the lag between trade-in and subsequent resale follows the geometric distribution from equation 3.

3.1 Data Consistency Tests

A major concern, given the low market share of the trade-ins data, is that the data are not equally representative over different game ages (time since release). If, for example, GameStop changes its used game acquisition policy with the age of the game such that the relative appeal of selling games via online auctions changes, then the appropriate factor for scaling up trade-ins in the data to represent the total trade-ins in the U.S. market may change as the games age. For example, the appropriate scale-up for the first two months may not be appropriate for, say, the sixth month.

The trade-ins data provide a simple check for this concern. Sellers whose shipments originated from zipcodes of the lowest population density quintile, average density 114 people per square mile, likely have poor access to brick and mortar stores like GameStop making offline trading-in impractical. Their propensity to resell via auction as opposed to a brick and mortar store should thus be unaffected by offer prices for used games trade-ins at brick and mortar stores. Consumers in high population density areas, however, should on average have much better access to brick and mortar stores to trade-in their used games. If the relative benefit of trading-in at GameStop change with games' ages, then consumers in high density areas would likely shift their trade-ins from nearby to online, or vice versa, as games age. Consumers in low density areas, however, likely resell used games online regardless. Therefore, if the appeal of trading-in at GameStop changes with a game's age, then the fraction of used game auctions for a game occurring in the first several months of release, relative to later following release, should be different for rural and urban consumers. Figure 1 shows this does not occur. The plot shows the average quantity of used game auction sales separately by the seller's zipcode's population density quintile, normalized to sum to 1 over the first year, separately for each quintile. No obvious differences are apparent between the pattern of used game auction sales originating in different population density quintiles, supporting the contention that the data are equally representative over game age, at least over the first year.

A second concern is that the data's market share of used game trade-ins changes over the time period analyzed, due to changing conditions or composition of the market. This too is easily checked. If this were to occur, then one would expect the ratio of trade-ins in the data to total new game sales in the U.S. market to change over time. Figure 2 shows no consistent trend of this sort, alleviating this concern as well.

A third concern is that the scale-up factor was not accurately estimated. To test the sensitivity to the assumptions used in constructing the scale-up, major analyses in later sections will be repeated assuming large differences in the scale-up factor for used game

auctions.

3.2 Summary Statistics

The trends in the new game price data, shown in Table 1, suggest that the decisions of when to buy and when to resell are non-trivial. If consumers buy a game right after it is released, they typically pay about \$55 for the game. But, if they wait to buy the game, they can acquire the game for much less, since the price typically declines rapidly. They can on average save 22% by waiting 6 months, and almost 50% by waiting a full year to buy the game. The implied rental prices (the buying price minus the amount received when reselling later) also typically decline over time. Since both prices and implied rental prices decline over time, consumers must trade off between buying and using the product immediately and paying less for it later.

The sales data, also summarized in Table 1, show that sales are more front-loaded than in the classic diffusion model, in which sales start slowly and increase over time. This front-loading may be due to the firm's response to competition from used sales. In the market for video games, on average about 40% of total sales of a game in the first year occur in the first two months. Not surprisingly, almost all of these games are new, and the firm profits directly from these sales. As time progresses, while total sales typically decline, the number of trade-ins initially increases, reaching a peak in the fifth month, and thereafter declines slightly before appearing to plateau. Used copy sales by retailers to consumers follow a similar pattern, with a slight lag, peaking in the 7th month following release. Since new sales continue to decline, the relative proportion of game sales that are of used copies increases over time. By the end of the first year, monthly used sales by retailers account for over 40% of per month sales of a game. Hence, later on, producers faces steep competition from used goods.

The cumulative sales data also suggest used game transactions are important in this market. By the end of the first year, 19% percent of new games sold have been traded-in.

The data exhibit obvious seasonality, e.g. more sales around Christmas. One method for controlling for seasonality is to add monthly demand shifters. However, Gowrisankaran and Rysman (2010) note the lack of intuition for why products would be enjoyed so much more during the Christmas season, also noting that adding season dummies to a dynamic model typically requires adding an additional state variable, substantially slowing estimation due to the curse of dimensionality. Gowrisankaran and Rysman (2010) show that deseasoning the data a priori yields similar results to traditional methods, suggesting a better alternative.

Each monthly numerical variable in the dataset is deseasoned by running a regression of its log on the composite critic review score and its square, age of game dummies, and date fixed effects. Specifically:

$$\begin{aligned} &Log(Dependent_t) = \\ &alpha + \beta_1 * rev_score_j + \beta_2 * rev_score_j\hat{}2 + \lambda_{age(t)} * I\left(age\left(t\right)\right) + \gamma_t * I\left(t\right) + \varepsilon_t \end{aligned} \tag{4}$$

The dependent variable is deseasoned by subtracting $\gamma_t * I\left(t\right)$ from the log of the dependent variable, and then exponentiating. This variable is then scaled up or down to the point where the sum of this seasonally adjusted variable equals the sum of the corresponding raw, (not deseasoned) variable.

3.3 Law of Motion

An important component of any dynamic model of consumer demand is the specification of consumer expectations of payoff relevant state variables in future periods. A common method is to assume that consumers base expectations on reduced form regressions of current state variables on lagged variables and product characteristics (Gowrisankaran and Rysman, 2012; Lee 2012b, 2010; Nair, 2007). The underlying assumption is that the reduced form model well replicates heuristics that consumers have developed with experience.

To see which observables predict future values of purchase prices, purchase prices were regressed on lags of price and game characteristics. The results are shown in Table 2. The table shows that lagged purchase price, critic review score quintiles, and month fixed effects are significant predictors. Twice lagged prices, genres, and other game characteristics, not shown, are not significant when added to the regression. The last column shows the results when adding the lagged residual from the prior column's regression as an explanatory variable. The t statistic on the corresponding coefficient tests the null assumption that the residuals are not autocorrelated (Wooldridge, 2002). It is not significant, suggesting autocorrelation is not a concern.

Table 3 repeats this exercise for trade-in prices of used games. The first column shows that the lagged purchase price is a reasonable predictor of the trade-in price consumers receive from selling used games. However, the lagged trade-in price is a better predictor. Including it as an explanatory variable, rather than lagged purchase price, raises the R^2 from 0.67 to 0.89. No other variable increases the R^2 by more than 0.01, suggesting that while other variables may be statistically significant, they are not meaningful to the consumer's decision. The autocorrelation test shows a similar result. Adding lagged residuals from the regression in column 6 yields a statistically significant coefficient estimate. However, it does not appear meaningful to consumers' decisions. It implies an effect of only -8 cents for each \$1 of residual the prior period.

4 Empirical Model

4.1 Demand

The model of consumer demand is cast as a discrete choice problem, where at the beginning of each period each consumer decides whether to be an owner of each game, independent of which other games they own. If, at the beginning of a time period, they do not own the product in question, this framework requires that they decide between buying and not buying the product that period. If they do own the product, they alternatively decide between keeping the product and selling it. If they buy or keep the product that period, they receive full flow utility from using the game. If they wait to buy or sell the product, they receive no flow utility from the game that period.

This setup implicitly makes several assumptions. First, it assumes consumers do not have use for more than one copy of a particular game. This seems reasonable given that the second copy does not provide any additional functionality, and since prices fall rapidly, there is little opportunity for arbitrage. Second, this framework implies that games are not substitutable for one another, and consumers do not explicitly choose between them. This assumption has been statistically verified on some platforms, but not all (see Nair, 2007; Derdenger, 2010). I repeat the statistical tests, and find no evidence that popular

games are substitutable. Details are provided in Appendix A. Although this assumption may seem counterintuitive to some, there are reasons to believe it roughly holds. Nair (2007) provides some reasons why video games are not likely substituted for each other. Another is that other entertainment goods such as movies might be just as substitutable, and with such a large number of potential substitutes, the maximum value of the marginal alternative might not vary much. This would imply that the assumption is a reasonable enough for the purposes of this exercise, which is consistent with the statistical tests. Third, this setup implies that used and new games are perfect substitutes for one another. Since used games provide the same service as new games, it seems quite reasonable that this assumption would nearly hold.

In the remainder of this section, I present the specifics of the demand model for a single game. I start by presenting the flow utilities. Next, I describe the transition processes of the state variables from the perspective of consumers. Then, I introduce the value functions. Finally, I describe the policy functions, which are an input in the estimation procedure. This process is repeated for each game during estimation. Game subscripts are omitted for parsimonious notation.

4.1.1 Flow Utility

There are four possible actions consumers can take at some point. They are buying, waiting to buy, keeping, and selling.

The flow utility of buying a game is given by:

$$u_{buy}\left(\bar{\delta}_i, \alpha_i, \xi, P, \varepsilon_i\right) = \bar{\delta}_i + \xi - \alpha_i P + \varepsilon_i = \bar{u}_{buy}\left(\bar{\delta}_i, \alpha_i, \xi, P\right) + \varepsilon_i \tag{5}$$

where $\bar{\delta}_i$ is the mean flow utility of the product to individual i at first use, ξ is a transient utility shock common across individuals, α_i is the price sensitivity of consumer i, P is the price at which the game can be bought, and ε_i is an individual specific shock. The function $\bar{u}_{buy}\left(\bar{\delta}_i, \alpha_i, \xi, P\right)$ equals the flow utility of buying minus ε_i, and will be used subsequently for notational purposes.

The mean flow utility of owning, assuming purchased in a previous period, is given by:

$$u_{own}\left(\bar{\delta}_i, \xi, h, \varepsilon_i\right) = \left(\bar{\delta}_i + \xi\right) B\left(h\right) + \varepsilon_i = \bar{u}_{own}\left(\bar{\delta}_i, \xi, h\right) + \varepsilon_i \tag{6}$$

where $B\left(h\right)$ is a function that reflects the decrease in value due to length of previous ownership h.

For several reasons, it is assumed that the rate at which consumers lose interest is the same across games. First, it seems unlikely that consumers anticipate the game-specific rate of boredom. Rather it is learned with experience as they play the game. Second, Game Informer review's "Replay Value" scores, which purportedly inform consumers of how quickly they will bore of the game, did not predict differences across games in the rate of price declines or the timing of sales. This suggests differences across games are not large. Third, while it was not chosen for this reason, this assumption has the added benefit of reducing the number of parameters to be estimated.

The function $B\left(h\right)$ is parameterized as $B\left(h\right) = \left(1 - \lambda_1\right)\exp\left(\lambda_2 h\right) - \lambda_1$, where $0 \leq \lambda_1 \leq 1, \lambda_2 \leq 0$. I assume that h is capped at $H = 12$, i.e. subsequent ownership periods beyond H do not change the value of the $B\left(\right)$ function. Due to identification concerns, individuals were restricted to have homogeneous rates of lost interest, i.e. the same $B()$ function. More flexible specifications were tested, such as combining an exponential and

linear functional form, relaxing the assumption that monthly value with use plateaus. However, none of the more flexible specifications meaningfully improved fit, or changed the shape of the estimated boredom function.

The mean flow utility of the outside good equals:

$$u_{wait}\left(\varepsilon_{i,0}\right) = \omega + \varepsilon_{i,0} = \bar{u}_{wait} + \varepsilon_{i,0} \tag{7}$$

where ω is the outside good's value and $\varepsilon_{i,0}$ is an individual specific shock. ω is normalized to a positive constant large enough to ensure that $\left(\bar{\delta}_i + \xi\right)$ is positive in every instance. As long as this holds, the actual value of ω is inconsequential. If and when $\left(\bar{\delta}_i + \xi\right)$ were less than zero, the model would imply tiring with use increases the valuation for it, which does not make sense.

The mean flow utility of selling is given by:

$$u_{sell}\left(P^{TI}, \alpha_i, \zeta, \varepsilon_{i,0}\right) = \omega + \alpha_i\left(P^{TI} + \zeta\right) + \varepsilon_{i,0} = \bar{u}_{sell}\left(P^{TI}, \alpha_i, \zeta\right) + \varepsilon_{i,0} \tag{8}$$

where P^{TI} is the price at which owners can resell the product, i.e. the trade-in price, and α_i is the same as in the buying equation. ζ, a transaction cost shock common across individuals, is similar to the demand shock ξ, but is experienced only by current owners.

4.1.2 Heterogeneity

Managers have noted the video game market consists of two groups of consumers, "hard-core gamers" and the "mass market."[22] I use a latent class approximation to the bimodal distribution of valuations (Kamakura and Russell, 1989). Following earlier papers focusing on the market for video games (Liu, 2010; Nair,2007) I assume there are two types. Type is denoted by k. The low type has intrinsic value of ownership equal to $\delta_k = \delta$, and the high type has intrinsic value equal to $\delta_k = \delta + \beta$. The price sensitivity is also estimated separately by type. The fraction of high types amongst the population is estimated as $\frac{exp(\gamma_1+\gamma_2*A_j^P)}{1+exp(\gamma_1+\gamma_2*A_j^P)}$, where A_j^P is the age of platform at the time of game j's release. The parameter γ_2 accounts for the changes in the initial composition of consumers, i.e. fraction high-type, in the market for the game as the platform matures. Lee (2012b) noted the importance of controlling for changes in the installed base of potential buyers in the context of video games.

The fraction of high types among non-owners and owners changes endogenously over time. In early periods, high type consumers are more likely to purchase, so naturally the fraction of remaining non-owners that are of the low type increases over time.

4.1.3 State and Control Space

While there is only one control variable, ownership, there are several other relevant state variables. They are the price (P), trade-in price (P^{TI}), previous periods of ownership (h), demand shock (ξ), transaction cost shock (ζ), and individual specific utility shocks (ε_i and $\varepsilon_{i,0}$). Ownership status and previous periods of ownership are deterministic. The remaining state variables are stochastic from the perspective of consumers.

Evidence presented in section 3.3 supports the assumption that purchase prices and trade-in prices follow a first order Markov process. Accordingly, I assume that, from the

[22] See Nair (2007).

perspective of consumers, purchase prices for the highest quintile of games follow the random process described in equation 9 below.

$$P_{t+1} = g\left(P_t\right) + \eta_{t+1} \tag{9}$$

where $g\left(P\right)$ is the function estimated from the data, and η is the component of the change in price unpredictable to consumers. Similarly, I assume consumer expectations of trade-in prices are given by:

$$P_{t+1}^{TI} = f\left(P_t^{TI}\right) + \eta_{t+1}^{TI} \tag{10}$$

where $f\left(P^{TI}\right)$ is the function estimated from the data, and η^{TI} is the component of trade-in prices unanticipated by consumers.

Following Nair (2007), I allow consumers to incorporate the correlation between price shocks η and demand shocks ξ in their expectations.[23] Specifically, I assume η and ξ are distributed jointly normal with correlation $\rho_{\xi,\eta}$ to be estimated. I also assume that these shocks are uncorrelated with other state variables. Since they are also not persistent, these shocks follow Rust's conditional independence assumption (Rust, 1987).

Each of the remaining stochastic state variables $(\zeta, \varepsilon_i, \varepsilon_{i,0})$ are assumed to be independently distributed across time, and hence are uncorrelated with all other state variables, present and future. This implies that they too follow Rust's (1987) conditional independence assumption. For functional forms, I assume ζ are distributed normally, and ε_i and $\varepsilon_{i,0}$ follow the type 1 extreme value distribution with location parameter equal to the negative of Euler's constant and scale parameter equal to one.

4.1.4 Value Functions

The ownership control variable is binary. Hence, for presentation purposes, the value function can be broken into two value functions, each conditional on ownership status.

For both value functions, one can yield a simplified Bellman equation on a reduced state space without the variables that satisfy Rust's conditional independence assumptions (ξ,ζ,η,η^{TI},ε_i, and $\varepsilon_{i,0}$) by integrating these states out of the value function. This results in the expected value of the value function before any of these variables are known. The states ε_i and $\varepsilon_{i,0}$ can be integrated over analytically, following Rust (1987). However ζ, ξ, η, and η^{TI}, must be integrated out numerically.[24]

I specify the "alternative specific" (i.e. choice specific) value functions so that they provide the expected maximum discounted future utility.[25] The equation for the "alternative specific" value function of owning (excluding current flow utility), equals:

[23] These assumptions imply that the firm can condition on the current value of the component of demand not observed by the econometrician, ξ, but not future ones, when setting prices.

[24] When ε_1 and ε_2 follow the type 1 extreme value distribution with location parameter equal to the negative of Euler's constant, and scale parameter equal to one: $E\left[\max\left(A+\varepsilon_1, B+\varepsilon_2\right)\right] = \ln\left(e^A + e^B\right)$. See Rust (1987), equation 4.12.

[25] Note that this is a nontraditional specification of the value function. It (1) does not include the current flow utility, and (2) is specified as the expected maximum utility from next period's options. Neither difference, however, precludes contraction - one can verify that this function satisfies the monotonicity and discounting sufficient conditions in Blackwell's Theorem. While one could specify the value function in the traditional way for this problem, doing so would slow estimation considerably because it requires ξ be included as a state variable in order to account for its correlation with η. This is because ξ cannot be integrated out without knowledge of η, which requires knowledge of P_{t-1} as well as P. Hence lagged price, or the price shock, would need to be included as a state variable too in order to integrate out ξ.

$$W_O(S_t) = \int \ln \left\{ \begin{array}{l} \exp\left(\bar{u}_{own}\left(S_{t+1}, \xi_{t+1}\right) + \varphi W_O\left(S_{t+1}\right)\right) \\ + \exp\left(\bar{u}_{sell}\left(S_{t+1}, \zeta_{t+1}\right) + \varphi \frac{\omega}{1-\varphi}\right) \end{array} \right\} df\left(S_{t+1}, \xi_{t+1}, \zeta_{t+1} | S_t\right) \quad (11)$$

where the state variables $S = \{\bar{\delta}_k, \alpha_i, P^{TI}, h\}$, φ is the discount factor, and $\varphi \frac{\omega}{1-\varphi}$ gives the expected discounted value of using the outside product for all future periods. This formula defines the expected future utility as the expected maximum utility from the consumer's two choices in the next period, assuming that the consumer owns the product at the end of the current period. If, next period, the consumer again keeps the product, they obtain flow utility from using it, plus the expected discounted utility from continuing to own it going forward. If, next period, they sell the product, they obtain utility from money received from selling it, plus the discounted expected utility of using the outside good for all future periods.

Likewise, the equation for the alternative specific value function for the future value of not having owned the product, W_{NO}, can be written as:

$$W_{NO}(S_t) = \int \ln \left\{ \begin{array}{l} \exp\left(\bar{u}_{buy}\left(S_{t+1}, \xi_{t+1}\right) + \varphi W_O\left(S_{t+1}\right)\right) \\ + \exp\left(\bar{u}_{wait} + \varphi W_{NO}\left(S_{t+1}\right)\right) \end{array} \right\} df\left(S_{t+1}, \xi_{t+1} | S_t\right) \quad (12)$$

where in this equation the state variables $S = \{\bar{\delta}_k, \alpha_i, P, P^{TI}, h\}$. This equation equals the expected maximum utility of the consumer's two choices in the next period, assuming the consumer does not own the product at the end of the current period. If, next period, the consumer continues to wait to buy, she receives the flow utility from using the outside good plus the expected discounted utility from continuing to not own the product. If, next period, the consumer buys the good, she receives flow utility from using the product minus the disutility of price plus the expected discounted utility of owning the product going forward. To reduce the computational burden of including several continuous state variables, I assume, for the above equation only, consumers use the purchase price to approximate the trade-in price according to column 1 of Table 3. Anecdotal evidence suggests consumers do not pay attention to the trade-in price when deciding whether to buy a game. Rather, consumers seem to use heuristics based on past experiences, supporting this assumption.

4.1.5 Policy Functions

An owner will sell the product if the expected discounted utility of selling exceeds the expected discounted utility of keeping. Specifically, the optimal policy is selling if and only if:

$$\bar{u}_{sell}(S, \zeta) + \varphi \frac{\omega}{1-\varphi} + \varepsilon_{i,0} > \bar{u}_{own}(S, \xi) + \varphi W_O(S) + \varepsilon_i \quad (13)$$

Since the error terms ε_i and $\varepsilon_{i,0}$ follow the type 1 extreme value distribution, the probability of selling for owner i of type k with h previous ownership periods can be written analytically as:

$$s_{k,sell}(S, \xi, \zeta) = \frac{\exp\left(\bar{u}_{sell}(S, \zeta) + \varphi \frac{\omega}{1-\varphi}\right)}{\exp\left(\bar{u}_{sell}(S, \zeta) + \varphi \frac{\omega}{1-\varphi}\right) + \exp\left(\bar{u}_{own}(S, \xi) + \varphi W_O(S)\right)} \quad (14)$$

Following analogous steps, the probability of buying for non-owner i of type k can be written as:

$$s_{k,buy}(S,\xi) = \frac{\exp\left(\bar{u}_{buy}(S,\xi) + \varphi W_O(S)\right)}{\exp\left(\bar{u}_{buy}(S,\xi) + \varphi W_O(S)\right) + \exp\left(\bar{u}_{wait} + \varphi W_{NO}(S)\right)} \tag{15}$$

5 Estimation

The model is estimated by maximum likelihood. The likelihood function equals:

$$L\left(P, P^{TI}, Q, Q^{TI}; \theta\right) = \prod_{j,t} L\left(P_{j,t}, P^{TI}_{j,t}, Q_{j,t}, Q^{TI}_{j,t}; \lambda, \alpha, \beta, \gamma, \delta_j, \sigma_\xi, \sigma_\eta, \sigma_{\eta^{TI}}, \sigma_\zeta, \rho_{\xi,\eta}, \varphi\right) \tag{16}$$

where j and t denote product and time, λ determines the rate of lost interest, α is the price sensitivity, β and γ determine the distribution of types, δ_j denotes game j's quality, σ_ξ, σ_η, $\sigma_{\eta^{TI}}$, and σ_ζ are the standard deviations of demand, prices, and transaction cost shocks, ξ, η, η^{TI} and ζ, $\rho_{\xi,\eta}$ is the correlation between ξ and η, and φ is the discount factor. Q is the quantity purchased, including used game purchases, and Q^{TI} is the quantity of trade-ins. Prior literature has noted problems in estimating the discount factor in dynamic models, and often assumes its value (Dube Hitch Chintagunta, 2010; Nair, 2007). Following Nair (2007), I assume the value of the discount factor equals 0.975. The discount factor is further discussed in section 6.2. The standard deviation of the price shocks, σ_η and $\sigma_{\eta^{TI}}$, can be estimated a priori.

After a change of variables transformation, the likelihood can be written in terms of the price, demand, and transaction cost shocks, as:

$$L(\eta, \xi, \zeta; \theta) = \prod_{j,t} f(\eta_{j,t}, \xi_{j,t}; \theta)\, f(\zeta_{j,t}; \theta)\, ||J|| \tag{17}$$

where $||J||$ is the Jacobian determinant. The derivation of the Jacobian is shown in Appendix B. The price shocks η are estimated a priori from the data, according to equation 8. The demand and transaction cost shocks (ξ and ζ) are recovered within each iteration, by the method described in the next subsection.

5.1 Recovering Error Terms

To compute the values of the demand and transaction cost shocks, ξ and ζ, for each product and period, I follow the method in Gowrisankaran and Rysman (2010), Nair (2007), and Schiraldi (2010). In each iteration in the maximization procedure, the value functions are computed first. Then, the shocks are computed as explained below.

For each game j, the values of ξ are found sequentially using Berry, Levinsohn, and Pakes' (1995) contraction mapping to find the value of $\xi_{j,t}$ which equalizes the observed and model's predicted aggregate share buying.[26] The formula for the model's predicted market share buying is:

[26] When ξ follows Rust's (1987) conditional independence assumption, the value function is independent of ξ, and the contraction mapping from Berry et al. (1995) still holds.

$$\bar{s}_{buy}\left(S_{j,t},\xi_{j,t}\right)=\frac{\sum\limits_{k}M_{j,k,t}*s_{k,buy}\left(S_{j,t},\xi_{j,t}\right)}{\sum\limits_{k}M_{j,k,t}} \tag{18}$$

where $M_{j,k,t}$ equals the mass of non-owners of type k in the market for game j at the beginning of period t.

The mass of non-owners in the first period equals the installed base of potential consumers in that period. In subsequent periods, it is updated to reflect buyers, who have exited the market, and incoming potential consumers that were not in the market for the game in the previous period, for example because they had not yet joined the platform, i.e. bought the console. The fraction of entering consumers of a given type then equals the total amount of entering consumers multiplied by the fraction of consumers that are of that type, implied by the γ parameters. The formula for updating the masses of non-owners in each period after the game is released is given below:

$$M_{j,k,t+1}=M_{j,k,t}*\left(1-s_{k,buy}\left(S_{j,t},\xi_{j,t}\right)\right)+M_{j,k}^{new} \tag{19}$$

where $M_{j,k}^{new}$ is the incoming group of potential consumers of type k entering each period.

The transaction cost shocks, ζ, can be found through a similar method, again separately for each game j. In each period beyond the first, $\zeta_{j,t}$ is found by equalizing the model's predicted and actual share of owners selling, using the contraction mapping in Berry, Levinsohn, and Pakes (1995). The formula for the predicted share selling is:

$$\bar{s}_{sell}\left(S_{j,t},\xi_{j,t},\zeta_{j,t}\right)=\frac{\sum\limits_{k,h}R_{h,j,k,t}*s_{k,sell}\left(S_{j,t},\xi_{j,t},\zeta_{j,t}\right)}{\sum\limits_{k,h}R_{h,j,k,t}} \tag{20}$$

where $R_{h,j,k,t}$ is the mass of owners of type k with h previous ownership periods of game j at the beginning of period t.

The mass of owners of each type and previous ownership length, $R_{h,j,k,t}$, evolves similarly to masses of non-owners. Before the product is introduced, there are no owners. After the good is released, the masses of owners are updated each period by:

$$R_{h,j,k,t+1}=\left\{\begin{array}{ll} M_{j,k,t}*s_{k,buy}\left(S_{j,t},\xi_{j,t}\right) & for\ h=1 \\ R_{h-1,j,k,t}*\left(1-s_{k,sell}\left(S_{j,t},\xi_{j,t},\zeta_{j,t}\right)\right) & for\ h>1 \end{array}\right\} \tag{21}$$

Hence the number of owners of game j of discrete type k with one period of previous ownership equals the mass of buyers of type k in the previous period. The mass of owners of type k with $h>1$ previous periods of ownership of game j equals the mass of individuals of type k, who in the previous period had $h-1$ previous ownership periods and decided not to sell.

5.2 Market Size

It is common in studies of video games (see for example Nair (2007) and Lee (2012b)) to assume that the market size is comprised of two components, initial consumers present when the game is released, and incoming consumers arriving each period. This is a natural assumption for two-sided markets, which requires the platform (hardware) to use software.

The easiest method for constructing market sizes, though not the best when modeling resale, is to assume all platform owners are in the market for a game. However, there are several reasons to believe that some platform owners may have no value for specific games, irrespective of quality, implying that they are not for practical purposes in the market for those games. For example, an adult male may have no interest in a game designed for young kids, even if it has exceedingly high quality. Moreover, the share of platform owners with even an iota of interest in a game may vary by game.

Previous research supports this assertion. Shiller and Waldfogel (2011) found using stated valuations that, for songs on iTunes top 50, on average about a third of individuals attribute zero value to the song. However, which individuals had zero valuations depended on which song. Moreover, the zero valuations could not reasonably be explained as being the censored tail of the distribution of valuations.

The video game data seem to display similar features. Some games, even among just those from the top quintile of critic review scores, sell to a low percentage of platform owners. Even among games from the top quintile of critic review scores, average shares of platform owners buying of a specific game vary greatly in the first year, from an about 0.8% to 55.3%, nearly a two orders of magnitude difference.

Falsely attributing differences in sales to quality levels could lead to odd results in models of resale markets. Under such assumptions, games with low sales would be estimated to have low quality. This would in turn imply that most purchases were due to transient utility shocks, and therefore owners would be keen on quickly reselling these games very soon after purchase. The model would have trouble explaining why they do not, and may lead to biased estimates of the rate of lost interest for games. Heterogeneity, depending on the form, may be able to absorb incorrect assumptions on market sizes common across games. But, it cannot explain differences across games unless heterogeneity is allowed to vary by game, which is not feasible - heterogeneity is not separately identified by game. Moreover, differing heterogeneity across games may merely be compensating for incorrect market size assumptions, suggesting a simpler, and feasible, approach.

The approach used in this paper uses multiple pieces of information to construct market sizes. It uses the number console owners to determine the overall magnitude of market sizes. It also uses the the number of sales after prices stabilize for a game to measure the relative market sizes across games, thus accounting for the empirical finding of Shiller and Waldfogel (2011). Logic supports this approach. Once the price has stabilized, there is no reason for consumers to delay purchase. At that point, any consumer with value above some threshold should buy the game. Specifics are explained below.

I employ a method related to the one in Nair (2007) to determine relative market sizes across games. For each game, I assume that the amount to which sales plateau after prices stabilize equals a set proportion κ of incoming potential consumers $M_{j,k}^{new}$. That is, I assume that size of incoming potential consumers is constant for a game. Hence, it can be observed from sales late in the product's lifecycle. I use only game's launched before the end of 2007, hence restricting the data to games observed for at least 24 months. I then find the number of such consumers in the market for the game when the product is launched by subtracting from total cumulative purchases in the last period the number of cumulative purchases due to consumers who entered the market after the product's launch. Last, I find the level of κ such that the game with the highest proportion of console owners consuming it has market size equal to the installed base of platform owners. Market sizes for all games are then scaled up from the two above components

using this value of κ.

5.3 Controlling for Endogeneity

A worry with this estimation approach is that the firms have more information on the demand shock ξ than the econometrician, which would give rise to a biased estimate of the price coefficient. Controlling for endogeneity is difficult in this context. A valid instrument would be a shock to costs or competitive conditions. It is difficult to find Bresnahan style instruments (Bresnahan 1981, 1987) for competitive conditions, since statistically games are not very substitutable (shown in Appendix A). It is also difficult to find strong Hausman style instruments for cost shocks (Hausman, 1996, Hausman and McFadden, 1984), since marginal production costs are very low (about $1.50).

The difficulty in finding valid instruments motivates use of an alternative method to control for endogeneity, from Villas-Boas and Winer (1999), Rossi et al. (2005), Nair (2007) and Jiang et al. (2009). The method involves first estimating a reduced form regression of price on an instrument for price, typically lagged price, and recording the residuals. In this context, residuals contain the endogenous component of price. Then, with these residuals known, the dependence of unanticipated price changes on demand changes can be explicitly accounted for by allowing the demand shocks and the price residuals to be correlated. The exact extent of the correlation is estimated within the model.

5.4 Identification

Heterogeneity in intrinsic valuations (β, γ) is identified by the average trends in prices and in the share of non-owners buying. To explain how, I borrow an argument from Nair (2007). Both the price and implicit rental cost, i.e. the price at time of purchase minus the amount received when selling the product, decline on average over time. However, the quality $\bar{\delta} + \xi$ on average stays the same. Hence, the expected gain from buying typically increases over time. If valuations were homogeneous, the share of non-owners buying should increase over time as well. If valuations are heterogeneous, higher valuation individuals buy with higher probability, leaving in later periods a greater share of low valuation types, who are less likely to buy than high types, implying the share buying can decline. The exact trends in share buying reflects the extent of heterogeneity. Differences across games released at different stages in the console's lifecycle identifies the impact of console's age on the composition of types, k.

The price sensitivities are identified by cross-sectional variation in prices and share of non-owners buying. Heterogeneity in this term is identified by the differential impacts of price shocks on sales early in the product's lifecycle, and later in the lifecycle when a greater share of non-owners are of the low valuation type and the low valuation type are the marginal buyers. This logic is similar to identification arguments from Lazarev (2012).

The coefficient of lost interest (λ) generates the time pattern of used sales over time. Higher boredom implies consumers are more likely to sell the product soon after purchase, which translates into a more active resale market.

The coefficient of lost interest and heterogeneity parameters are separately identified. While an infinite number of sets of rates of lost interest and initial valuations can result in an individual wanting to sell a game after h periods but not after $h-1$ periods, given

prices, they cannot explain the individual's choice of when to buy. Her first h uses must in expectation be valued at least high enough to justify buying when she did, but not high enough to justify buying earlier when prices were higher. A series of inequalities like these separately identifies these parameters.

5.5 Counterfactuals' Computation

Counterfactual simulations using the parameters from demand estimation allow comparison of prices, profits, and welfare under alternative environments. Specifically, I compare the status quo environment where resale is allowed with hypothetical environments where the firm sells only non-resellable goods, or where the firm exclusively rents goods directly to consumers. I do not consider the case of combining selling and renting.

Calculation of dynamic equilibrium prices in counterfactual simulations requires a different assumption on the evolution of prices than in demand estimation. In demand estimation, because the observed price path was an equilibrium outcome under existing laws, the current value of price characterized consumer expectations of future prices, and was therefore included as a state variable. This same price path cannot be assumed to be an equilibrium in a counterfactual environment. Instead, the prices, which are set optimally by the firm each period to maximize discounted profits, should depend on market conditions. Demand conditions for non-resellable goods can be characterized by the mass of non-owners of the good by consumer type (two variables). When resale is allowed, however, conditions for the supply of used goods, which compete with new goods, are relevant for pricing. Thus the masses of owners by type and ownership length, 24 variables in total, must also be included as continuous state variables. The same is true for renting, because demand to rent depends on the length of time different consumers have used the product.

While a method from Nair (2007) can be used the calculate the Markov perfect equilibrium prices when resale is prohibited, the large number of variables characterizing market conditions in the other cases, which need to be included as state variables, pose a challenge for estimation. I circumvent this problem for the case of allowed resale by taking observed prices and quantities sold from the data. By doing so, I am implicitly assuming that the observed prices would result from an equilibrium simulated by the model, were such a simulation feasible. For the case of rentals, I make some strong assumptions to allow simulation. However, since each of these assumptions are chosen to bias profits under renting upwards, it is still possible to calculate a lower bound for the amount by which selling non-resellable goods raises profits over renting.

5.5.1 Non-resellable goods

To find the price paths arising from a Markov Perfect rational expectations equilibrium game between firms and consumers when resale markets are shut down, I use a policy function iteration procedure similar to Nair's (2007).[27] The resulting policy functions for each player account for their own impact on the evolution of states, and are optimal given the calculated policy functions of the other players in the game.

The profit-maximizing firm policy function follows from the static profit function and the evolution of states. The static profit function is:

[27] Nair notes that the the equilibrium is not guaranteed to be unique. However, we both found that multiple starting values converged to the same equilibrium. This does not rule out a boundary solution.

$$\pi\left(S_{j,t}, P_{j,t}\right) = \left(P_{j,t} - MC\right) * Q\left(S_{j,t}, P_{j,t}\right) \tag{22}$$

where the state variables S include the masses of non-owners by type. $Q\left(S, P\right)$, the quantity purchased, is given by the product of the shares buying, described in equation 15, and the market sizes of each type. The firm's value function equals:

$$V_{Firm}\left(S_{j,t}\right) = \max_{P_{j,t}} \left\{\pi\left(S_{j,t}, P_{j,t}\right) + \varphi V_{Firm}\left(S_{j,t+1} | S_{j,t}, P_{j,t}\right)\right\} \tag{23}$$

where here the set of state variables, S, are the masses of non-owners of each type. These variables evolve according to equation 19. The firm's policy function is simply the profit maximizing price at each state:

$$P\left(S_{j,t}\right) = \arg\max_{P_{j,t}} \left\{\pi\left(S_{j,t}, P_{j,t}\right) + \varphi V_{Firm}\left(S_{j,t+1} | S_{j,t}, P_{j,t}\right)\right\} \tag{24}$$

The numerical procedure proceeds by four steps, separately for each game. The first step is simply to guess at both the consumers' policy functions, $s_{k,buy}\left(S\right)$, and their expectations for the evolution of the state variables, $M_{j,t+1}\left(s_{k,buy}\left(S\right)\right)$. The second step calculates an updated guess at the the firm's value and policy functions, $V_{Firm}\left(S\right)$ and $P\left(S\right)$, conditioning on the most recent guess of the consumers' policy functions, $s_{k,buy}\left(S\right)$. The third step iterates on two substeps. The first substep is calculating the next period states, $M_{j,t+1}\left(s_{k,buy}\left(S\right)\right)$, as a function of current states and the most recent guess at the consumer's policy functions, $s_{k,buy}\left(S\right)$. Then in the next substep the consumers' policy functions, $s_{k,buy}\left(S\right)$, are recomputed using the updated guess of consumers expectations, $M_{j,t+1}\left(s_{k,buy}\left(S\right)\right)$. These two substeps are repeated until the total absolute difference in the consumers' policy functions between iterations on the substeps is sufficiently small. This concludes the third step, yielding the next guess of the consumers' policy functions, $s_{k,buy}\left(S\right)$. The fourth step iterates on steps two and three repeatedly until the total absolute difference in the firm's policy function betweens iterations is smaller than some arbitrary value, at which time the procedure stops.

5.5.2 Rented goods

Many information goods are now "'streamed"', whereby the firm sells the right to use the digital product over the internet for a set length of time, without granting ownership, rather than outright selling the product. In this case, consumers lose access to the product at the end of the period, and must rent the product again if they wish to continue to have access to it. Hence, this strategy prevents resale. This strategy is commonly employed for movies (e.g. Netflix, Amazon Prime) and music (e.g. Spotify).

Several complications arise when attempting to estimate revenues. To circumvent these issues, four different assumptions are chosen for the rental simulations, each of which should bias rental profits upwards, yielding an upper bound estimate of rental profits. First, fully characterizing demand with state variables, necessary to compute a Markov perfect equilibrium, would require including a state variable for each combination of consumer type and length of ownership, too many to feasibly use in simulation. I therefore allow the firm to set flat rental prices. Second, the licensing fees the platform charges content producers, currently charged per physical copy produced, would likely be different for digital rentals. To address this, I assume zero marginal costs of renting, but assume an additional fixed cost equal to the total licensing cost incurred when resale

is allowed. Third, I assume zero transactions costs for returning or acquiring a rental and allow consumers to rent again after they had stopped renting. Fourth, consumers incorporate expected future value into their willingness to pay when purchasing a game, implying profits from outright sales during any length of time reflect consumer's use beyond that time. This would not be true, however, for short rentals. I therefore allow the firm to earn profits from renting over a much longer period when comparing profits, two years as opposed to one. However, to be consistent, I only allow new consumers to enter during the first year following release.

Optimal profits are simulated by searching over the flat rental price to maximize profits. For each rental price, the consumer policy functions are calculated. Their calculation follows steps similar to the steps in earlier sections.

6 Results

The parameter estimates and their standard errors are reported in Table 4. However, the meaning does not directly follow from their reported values. Below, I provide context to the estimates.

The rate at which consumers tire of products and the heterogeneity in initial valuations are depicted graphically in Figure 3. The Y-axis denotes dollar value consumers obtain for one month's use of a game. The X-axis denotes how many periods the individual has already owned the product. Note that this is not the age of the game, but rather how many periods a particular individual has owned the game. The two curves, one for each consumer type, show the consumers' average monthly values for owning a game conditional on how many periods they have owned that game previously. The "hardcore gamers," i.e. high types, value a month of ownership of an average game at about $80 if they have not owned it previously, about $30 more than the low types.

As consumers use and grow tired of a game, both types substantially lower their valuations for continued use of it. The high types lower their value for a month's use of an average game from about $80 for the first month they own it to about $4 per month after having owned it for 6 months. Moreover, about 90% of this decline occurs in the first two months following purchase.

The fraction of consumers in the market that have high intrinsic value for a game, as opposed to low intrinsic value, is implied by the parameters γ. For games released at the same time as the platform was released, it is estimated that about half of the consumers in the market for the game are of the high valuation type. For games released 6 months later, however, only about 20% are. This reflects the endogenous adoption decision shown in Lee (2012b) and found in Liu (2010). Consumers with high valuations for games presumably are more likely to buy the console early on, even before many games have been developed for it. Later, lower valuation types, often referred to in the video game industry as the "mass market," comprise a larger share of consumers.

The price elasticities over time are shown in Table 5. Static (one-period) elasticities and dynamic elasticities are reported. For both types of elasticities, I allow expectations of future prices to change with the price change. The average of the static elasticities, -2.16, is quite close to the average elasticity estimated in Nair (2011), which was also about -2. The dynamic elasticities, which report the percent change in the next 3 months sales for a current one percent change in price in one month, are similar.

The estimated values of the standard deviations of ξ and ζ are both relatively small.

The standard deviation of the demand shock, ξ, equals about \$4 in dollar terms for the high type, roughly 5% of the high type's typical initial one-period valuation for the good. The standard deviation of the transaction cost shock, ζ, is \$15 for the high type, roughly 15% of the high type's typical initial static valuation for the good. The correlation between the demand and price shocks is estimated to be 0.24.

6.1 Counterfactual Results

The main findings are shown in Table 6. Profits from selling non-resellable games in the first year following a game's release are about twice as high as the corresponding profits from selling resellable goods. This is true before accounting for the fixed costs of developing the game.[28] There is substantial variation across games. The standard deviation of the percent increase in profits is 75%, although each simulated game is estimated to yield higher profits under prohibited resale. The overall difference is due to the fact that the firm sells more goods when resale is allowed, and that the goods sell for higher prices in later periods. Some of the difference in quantities sold is due to new sales being displaced by used copies when resale is allowed and some to the fact that consumers delay purchase beyond a year when resale is allowed and prices fall quickly. These latter sales are not captured in this comparison. Most individuals buying later when resale is allowed, however, buy used games if available, or buy the game at a price close to marginal costs, so the firm earns little profits on later sales when resale is allowed. Hence ignoring later sales has negligible impact.[29]

The rental simulations indicate the upper bound estimate of profits is barely higher than profits from selling a resellable good, and still much lower than profits from selling a non-resellable good. Profits from selling a non-resellable good are still roughly double the profits from renting.

The estimated profits from rentals provided credibility to the overall findings. The standard approach for comparing profits between the status quo and counterfactual policies is to simulate both scenarios using the estimated parameters. For rentals and non-resellable goods, this was feasible. However, for reasons described earlier, it is not feasible to fully simulate the case where resale is allowed. This leads to the concern that a poorly specified model might be driving the estimated increase in profits from preventing resale. However, the observed relationship between rental profits and resellable good profits can be used to address concerns. Because buying and selling a resellable good is logically similar to renting it, but the implied rental prices with resellable goods are less flexible than explicit rental prices, we might expect profits from renting to be weakly higher than profits from resellable goods. This provides a check - are rental profits weakly higher than profits from selling a resellable good? They are, consistent with expectations.

The average price paths under each distribution strategy are shown in Figure 4. When resale is prohibited, optimal prices decline very slowly over time, falling by a few dollars

[28] Average development costs for XBOX360/PS3/PC based games for one major producer, Ubisoft, are found to average somewhere between \$18.8 million and \$28.2 million. Note that these costs can be recouped, however, from sales across multiple platforms and countries, and the profit figures in this paper only include US sales on the XBOX 360 platform. (http://www.gamasutra.com/php-bin/news_index.php?story=18389, accessed Nov 28, 2011.)

[29] For the 17 games released in the first two months of the dataset, 92% of profits (in the first three years) occur in the first year, on average. Assuming that this holds for all games, then profits from resellable goods are underestimated by about 8%, - not enough to explain the difference in profits when resale is shut down.

over the first year. As a result, consumers do not have much incentive to wait. However, when resale is allowed, price falls sharply over time. The price declines for this group of games by about \$30 over the first year on average, and implied rental prices fall too.[30] Consumers in the model anticipate the falling price, and fewer are willing to pay the high price initially, despite the fact that the good is worth more to them when resale is allowed, since they then have the option to resell the good.

The above finding is quite logical. When resale is allowed, lowering price reduces the number of traded-in games. Fewer trade-ins in the current period imply lower inventory levels of used games at retailers, which in turn implies that in subsequent periods fewer new copy sales will be displaced by used copy sales. Hence, by dropping price in the current period the seller reduces competition from used copies in later periods. This dynamic effect provides additional motivation to lower prices when resale is allowed.

A related question is what impact resale markets have on consumer welfare, after accounting for the firm's response. I calculate the exact net consumer welfare gain from resale markets using Small and Rosen's (1981) formula for Hick's equivalent variation. The net consumer welfare gain/loss from allowing resale to an individual of type k, in dollars, equals:

$$EV_k = \frac{W_{NO,YR}(k,1) - W_{NO,NR}(k,1)}{\alpha_k} \tag{25}$$

where α_k is the price sensitivity, and $W_{NO,YR}(k,1)$ and $W_{NO,NR}(k,1)$ are the value functions at game release, conditional on not owning the good, when resale is and is not allowed, respectively. I find that only distributing non-resellable goods would lower welfare by \$23.80 on average for each low-type individual, but raise welfare by \$11.87 on average per high type individual, for an aggregate loss of \$21 million on average across games.[31] High type individuals are better off because they can acquire the game for less. Low types, however, yield positive utility from use of the product for fewer periods. Losing the ability to resell the game is therefore quite costly to them.

Using an analogous equation for consumer welfare changes, I find that renting as opposed to selling resellable goods lowers aggregate consumer welfare by \$75 million for the average game. This loss arises from the fact that consumers forgo using the product when their valuations are between zero and the rental price.

Ishihara and Ching (2012) find qualitatively similar results for the impact of shutting resale in Japan. They find less of a difference in profits, possibly due to differences across these markets. In Japan, new games prices do not fall over time. This stands in stark contrast to other countries. Additionally, the Japanese market seems to highly value novelty. Accordingly, Ishihara and Ching (2012) allow for perceived quality to decline to non-owners as well, though they do not allow for heterogeneous usage values. [32] The fact that they find qualitatively similar findings despite investigating a market with obvious differences, and using alternative assumptions, supports the findings in both papers.

[30] There are multiple possible definitions for implied rental prices. If it is defined by the cost of buying minus the amount received from selling the product one period later, rental prices decline by about 80% over the first year.

[31] The welfare calculation does not include a third group, video game sellers, whose profit would increase welfare when resale is allowed. Their profits cannot be estimated because their cost functions are not known.

[32] Heterogeneity in usage values are not separately identified from novelty effects, restricting both from being estimated.

6.2 Robustness Checks

This subsection tests whether any of several modeling assumptions drives the main result. One at a time, each assumption is changed, and the estimation model and counterfactual simulations are rerun. The results are shown in Table 7.

The first concern is that any imprecision in the process of scaling up the secondhand sales data to proxy for the entire market might strongly impact the results. A robustness check shows that this does not happen, at least for small levels of imprecision. Comparing the first and second columns of Table 7 shows that a 25% reduction in the number of games traded-in increases estimated rental profits by 5%, and reduces estimated non-resellable good profits by 3%. By construction, profits from resellable goods are unchanged.

The second concern is that the assumed degree of forward-lookingness could have a large impact on the results. The monthly discount factor was assumed to equal 0.975, which corresponds to a yearly discount factor of 0.74, the same value as in Nair's (2007) paper. Although suitable variation did not exist to estimate the discount factor within the model, the assumed value falls between others' estimates. Lee (2012) estimated a monthly discount factor of 0.93, whereas Chevalier and Goolsbee (2010) found little evidence of any discounting, implying a discount factor close to 1. As a robustness check, I re-estimate profits under each scenario assuming the monthly discount factor equals 0.95. Results are shown in column 3. Non-resellable good profits are 4% lower, rental profits 29% lower, and resellable good profits 4% lower than in the main specification, suggesting small errors in the assumed value do not change the conclusions dramatically.

The third robustness check tests the importance of the assumed market size magnitude. The main specification assumed a market-size magnitude such that the most popular game had a market size equal to the installed base of platform owners. It is possible, however, that the true market size could be smaller. Reducing the market size by 25% demonstrates that the market size does influence the results. Non-resellable good profits were 25% lower, and rental profits 29% lower. This is likely due to the fact that increasing the market size increases the number of unserved consumers in the status quo case, who may be served and thus contribute to revenues in counterfactual environments.

This observed dependency motivated running another more extreme robustness check, reducing the market size by 50%. Under this assumption, nearly everyone in the market for a game buys it within the first two years. Lowering the market size much further results in some cases where the share non-owners purchasing exceeds one, causing the estimation procedure to fail. With a 50% reduced market size, profits from non-resellable goods are estimated to be 2% less than resellable good profits. However, to reach this qualitatively different result, a drastically different market size assumption was required.

Throughout all robustness checks profits from non-resellable goods were estimated to be much higher than rental profits. The percent increase ranged from 83% to 165%.

Because the reasoning behind why rental profits may be lower than non-resellable good profits extends to resellable goods, we might expect their profits, i.e. of rentals and resellable goods, to be similar. The main specification found that in fact they were: 522 vs. 495 million. In the robustness checks with different market size assumptions, their profits diverged. These results thus provide some ex-post support for market size magnitude assumed in the base model.

7 Discussion and Conclusion

Firms may soon be able to circumvent the first sale doctrine and legally prevent resale outright in the U.S. by selling non-resellable digital downloads or by streaming rentals. The results in this paper show that profits from selling video games would be much higher for non-resellable goods, over 100% higher in the base specification. However, the exact amount increase was sensitive to the assumed market size. Throughout all robustness checks, non-resellable good profits were substantialy higher than rental profits, around 100% on average.

These findings stand in contrast to most previous empirical and theoretical papers focusing on products for which consumers do not lose interest. They generally found monopolist's profits were not reduced by frictionless resale. The different results, along with the logic in section 2.1, suggest the rate at which consumers lose interest might be important in determining the impact of resale on producer profits. This may interest managers and policy makers, since the types of goods matching this assumption, such as entertainment information products (e.g. books), are the types of goods easily distributed as non-resellable downloads.

A Evidence that Games are Not Substitutes

Nair (2007) argued that because "each game is fairly unique, having its own distinct features, characters and idiosyncrasies," that they are likely not substitutes for one another. Empirical results showed (1) that there are no statistically significant cross-price effects within genre, (2) neither sales nor prices are significantly impacted by hit game releases in the same genre, and (3) that concentration in a genre does not significantly impact the rate at which prices decline. However, Derdenger (2010) yielded opposing results when investigating other 6th generation consoles.

To test whether XBOX 360 video games in my data are substitutes, I employ a static nested logit regression similar Nair's (2007) and Derdenger's (2010), where genre determines the nest. Following their specification, I regress:

$$\ln\left(s_{jt}/s_{0t}\right) = \alpha_j + \omega_t + \lambda\left(t - r_j\right) + \beta P_{jt} + \sigma \ln\left(s_{jt|g}\right) + \varphi_{jt} \tag{26}$$

where s_{jt} and s_{0t} are the shares of game j and the outside good in period t, respectively, α_j and ω_t are game and month fixed effects, r_j is the release date of game j, P_{jt} is the price, $s_{jt|g}$ is game j's share of genre g sales, and $\varphi_{j,t}$ is an error term. Following earlier papers (Nair, 2007; Derdenger, 2010), I set the market size to the cumulative installed base of XBOX 360 console owners. The value of the parameter σ determines the extent to which games are substitutes for one another. A value close to zero, rather than a higher positive number close to 1, denotes that the nests are largely irrelevant and hence games are not substitutes for one another.

Two right-hand size variables in equation 31 above, price (P_{jt}) and the within genre share of sales ($s_{jt|g}$), are likely endogenous. To address this concern, I instrument for them using the lagged price and the number of games in the genre.

The results are shown in Table 8. The sample is restricted to games with the highest quintile of critic scores. The first three columns in Table 8 employ OLS, and the second three employ $2SLS$. Note that once the instruments are included, the estimated values

of σ are negative and insignificant, strongly suggesting that games are not substitutable for one another.

B Jacobian

Note, the quantity predicted to be bought, Q^{pred}_{bought}, equals:

$$Q^{pred}_{bought}(S,\xi,M) = \sum_k M_k s_{k,buy}(S,\xi) \tag{27}$$

Similarly, the quantity predicted to be traded-in in a period, Q^{pred}_{resold} equals:

$$Q^{pred}_{resold}(S,\xi,\zeta,R) = \sum_{k,h} R_{k,h} s_{k,sell}(S,\xi,\zeta) \tag{28}$$

From these equations, and the pricing equation, the shocks (ξ, η, and ζ) can be found. ξ is found by solving:

$$Q^{pred}_{bought}(S,\xi,M) - Q^{actual}_{bought} = 0 \tag{29}$$

and,

$$\xi = Q^{-1}_{bought}\left(S, M, Q^{actual}_{bought}\right) \tag{30}$$

η is found by:

$$\eta_{t+1} = P_{t+1} - \kappa P_t \tag{31}$$

and ζ is found by solving:

$$Q^{pred}_{resold}(S,\xi,\zeta,R) - Q^{actual}_{resold} = 0 \tag{32}$$

$$\Rightarrow \zeta = Q^{-1}_{resold}\left(S,\xi,R,Q^{actual}_{resold}\right) \tag{33}$$

The Jacobian determinant equals:

$$\begin{vmatrix} \frac{\partial\xi}{\partial Q^{actual}_{bought}} & \frac{\partial\xi}{\partial P} & \frac{\partial\xi}{\partial Q^{actual}_{resold}} \\ \frac{\partial\eta}{\partial Q^{actual}_{bought}} & \frac{\partial\eta}{\partial P} & \frac{\partial\eta}{\partial Q^{actual}_{resold}} \\ \frac{\partial\zeta}{\partial Q^{actual}_{bought}} & \frac{\partial\zeta}{\partial P} & \frac{\partial\zeta}{\partial Q^{actual}_{resold}} \end{vmatrix} = \begin{vmatrix} \frac{\partial\xi}{Q^{actual}_{bought}} & \frac{\partial\xi}{\partial P} & 0 \\ 0 & 1 & 0 \\ 0 & \frac{\partial\zeta}{\partial P} & \frac{\partial\zeta}{\partial Q^{actual}_{resold}} \end{vmatrix} = \left| \frac{\partial\xi}{\partial Q^{actual}_{bought}} * \frac{\partial\zeta}{\partial Q^{actual}_{resold}} \right| \tag{34}$$

B.1 Computing $\frac{\partial\zeta_t}{\partial q^{resold}_t}$

We know that $\sum_{k,h} R_{k,h} s_{k,sell}(S,\xi,\zeta) - Q^{actual}_{resold} = 0$. By the implicit function theorem, we have:

$$\frac{\partial\zeta}{\partial Q^{actual}_{resold}} = -\frac{-1}{\frac{\partial}{\partial\zeta}\left(\sum_{k,h} R_{k,h} * s^k_{sell}(S,\zeta)\right)} \tag{35}$$

Note

$$\frac{d}{dx}\frac{e^{A+x}}{e^{A+x}+e^{B}} = e^{B}\frac{e^{A+x}}{\left(e^{B}+e^{A+x}\right)^{2}} = \frac{e^{B}}{\left(e^{A+x}+e^{B}\right)}\frac{e^{A+x}}{\left(e^{A+x}+e^{B}\right)} \tag{36}$$

Hence,

$$\frac{\partial}{\partial\zeta}s_{k,sell}\left(S,\xi,\zeta\right) = \left(1-s_{k,sell}\left(S,\xi,\zeta\right)\right) * s_{k,sell}\left(S,\xi,\zeta\right) \tag{37}$$

and we can rewrite:

$$\frac{\partial\zeta}{\partial Q^{actual}_{resold}} = -\frac{-1}{\sum\limits_{k,h} R_{k,h}\left(1-s_{k,sell}\left(S,\xi,\zeta\right)\right) * s_{k,sell}\left(S,\xi,\zeta\right)} \tag{38}$$

B.2 Computing $\frac{\partial\xi_t}{\partial q_t^{bought}}$

We know that: $\sum\limits_{k} M_k s_{k,buy}\left(S,\xi\right) - Q^{actual}_{bought} = 0$. By the implicit function theorem, we have:

$$\frac{\partial\xi_t}{\partial Q^{actual}_{bought}} = -\frac{-1}{\frac{\partial}{\partial\xi}\left(\sum\limits_{k} M_{k,t}s_{k,buy}\left(S,\xi\right)\right)} \tag{39}$$

Through the same steps, we can find that this equals:

$$\frac{\partial\xi}{\partial Q^{actual}_{bought}} = -\frac{-1}{\sum\limits_{k} M_k\left(1-s_{k,buy}\left(S,\xi\right)\right) * s_{k,buy}\left(S,\xi\right)} \tag{40}$$

References

[1] Anderson, S.P., Ginsburgh, V.A.: Price discrimination via second-hand markets. European Economic Review **38**(1), 23–44 (1994). URL http://ideas.repec.org/a/eee/eecrev/v38y1994i1p23-44.html

[2] Benjamin, D.K., Kormendi, R.C.: The interrelationship between markets for new and used durable goods. Journal of Law and Economics **17**(2), 381–401 (1974). URL http://ideas.repec.org/a/ucp/jlawec/v17y1974i2p381-401.html

[3] Berry, S., Levinsohn, J., Pakes, A.: Automobile prices in market equilibrium. Econometrica **63**(4), 841–890 (1995)

[4] Bresnahan, T.: Departures from marginal-cost pricing in the american automobile industry. Journal of Econometrics **17**(2), 201–227 (1981)

[5] Bresnahan, T.: Competition and collusion in the american automobile oligopoly: The 1955 price war. Journal of Industrial Economics **35**(4), 457–482 (1987)

[6] Bulow, J.I.: Durable-goods monopolists. Journal of Political Economy **90**(2), 314–332 (1982)

[7] Chen, J., Esteban, S., Shum, M.: How much competition is a secondary market. Working Paper (2011)

[8] Derdenger, T.: Technological tying and the intensity of competition: Empirical analysis of the video game industry. Working Paper (2011)

[9] Dube, J.P., Hitsch, G., Chintagunta, P.: Tipping and concentration in markets with indirect network effects. Marketing Science **29**(2), 216–249 (2010)

[10] Einav, L.: Seasonality in the u.s. motion picture industry. Rand Journal of Economics **38**(1), 127–145 (2007)

[11] Gowrisankaran, G., Rysman, M.: Dynamics of consumer demand for new durable goods. NBER Working Paper (2009). URL http://ideas.repec.org/p/nbr/nberwo/14737.html

[12] Graham, J.: Preserving the aftermarket in copyrighted works: Adapting the first sale doctrine to the emerging technological landscape. Stanford Technology Law Review **1**, 1–4 (2002)

[13] Hausman, J.: Valuation of New Goods under Perfect and Imperfect Competition, vol. 58 (1996)

[14] Hausman, J., McFadden, D.: Specification tests for the multinominal logit model. Econometrica **52**(5), 1219–1240 (1984)

[15] Hendel, I., Lizzeri, A.: Interfering with secondary markets. RAND Journal of Economics **30**(1), 1–21 (1999). URL http://ideas.repec.org/a/rje/randje/v30y1999ispringp1-21.html

[16] Hinkes, E.M.: Access controls in the digital era and the fair use/first sale doctrines. Santa Clara Computer and High Technology Law Journal **23**, 685–726 (2006)

[17] Hurt, H.: Turning whimsy into a gaming classic maybe. The New York Times p. 6 (2009)

[18] Ishihara, M., Ching, A.: Dynamic demand for new and used durable goods without physical depreciation: The case of japanese video games. Working Paper (2011)

[19] Jiang, R., Manchanda, P., Rossi, P.: Bayesian analysis of random coefficient logit models using aggregate data. Journal of Econometrics **149**(2), 136–148 (2009)

[20] Kamakura, W., Russell, G.: A probabilistic choice model for market segmentation and elasticity structure. Journal of Marketing Research **26**(4), 379–390 (1989)

[21] Lazarev, J.: The welfare effects of intertemporal price discrimination: An empirical analysis of airline pricing in u.s. monopoly markets. Working Paper (2012)

[22] Lee, R.S.: Home video game platforms. The Oxford Handbook of the Digital Economy (2012)

[23] Lee, R.S.: Vertical integration and exclusivity in platform and two-sided markets. Working Paper (2012)

[24] Leslie, P., Sorensen, A.: The welfare effects of ticket resale. Working Paper (2011)

[25] Liu, H.: Dynamics of pricing in the video game console market: Skimming or penetration. Journal of Marketing Research **47**(3), 428–443 (2010)

[26] Long, H.S.: Reconsidering the balance of the digital first sale debate: Re-examining the case for a statutory digital first sale doctrine to facilitate second-hand digital media markets. Alabama Law Review **59**(4), 1183–1202 (2007)

[27] Miller, H.L.: On killing off the market for used textbooks and the relationship between markets for new and secondhand goods. Journal of Political Economy **82**(3), 612–619 (1974). URL http://ideas.repec.org/a/ucp/jpolec/v82y1974i3p612-19.html

[28] Nair, H.: Intertemporal price discrimination with forward-looking consumers: Application to the us market for console video-games. Quantitative Marketing and Economics **5**(3), 239–292 (2007). URL http://ideas.repec.org/a/kap/qmktec/v5y2007i3p239-292.html

[29] Rich, M.: Declining book sales cast gloom at an expo. The New York Times p. 4 (2009)

[30] Rossi, P., Allenby, G., McCulloch, R.: Bayesian Statistics and Marketing. Wiley Series in Probability and Statistics (2005)

[31] Rust, J.: When is it optimal to kill off the market for used durable goods? Econometrica **54**(1), 65–86 (1986). URL http://ideas.repec.org/a/ecm/emetrp/v54y1986i1p65-86.html

[32] Rust, J.: Optimal replacement of gmc bus engines: An empirical model of harold zurcher. Econometrica **55**(5), 999–1033 (1987)

[33] Schiraldi, P.: Automobile replacement: a dynamic structural approach. Rand Journal of Economics **42**(2), 266–291 (2011)

[34] Seringhaus, M.: E-book transactions: Amazon Kindles the copy ownership debate. Yale Journal of Law and Technology **12**, 147 (2009)

[35] Shiller, B., Waldfogel, J.: Music for a song: An empirical look at uniform song pricing and its alternatives. Journal of Industrial Economics **59**(4), 630–660 (2011)

[36] Small, K., Rosen, H.: Applied welfare economics with discrete choice models. Econometrica **49**(1), 105–130 (1981). URL http://ideas.repec.org/a/kap/qmktec/v5y2007i3p239-292.html

[37] Villas-Boas, J.M., Winer, R.S.: Endogeneity in brand choice models. Management Science **45**(10), 1324–1338 (1999)

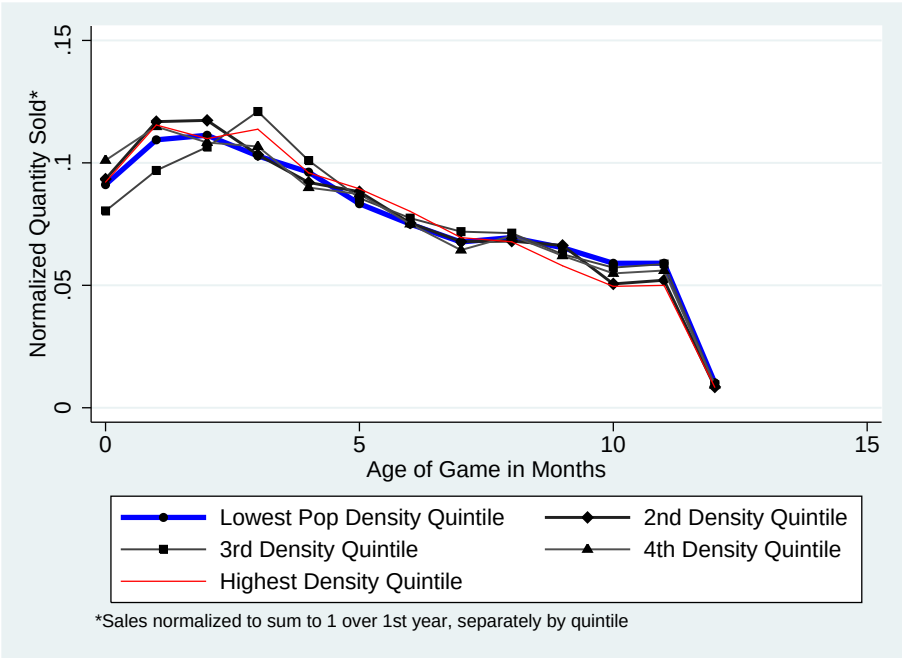

Figure 1: Normalized Used Auction Sales vs. Time Since Release - By Population Density Quintile of Seller's Zipcode

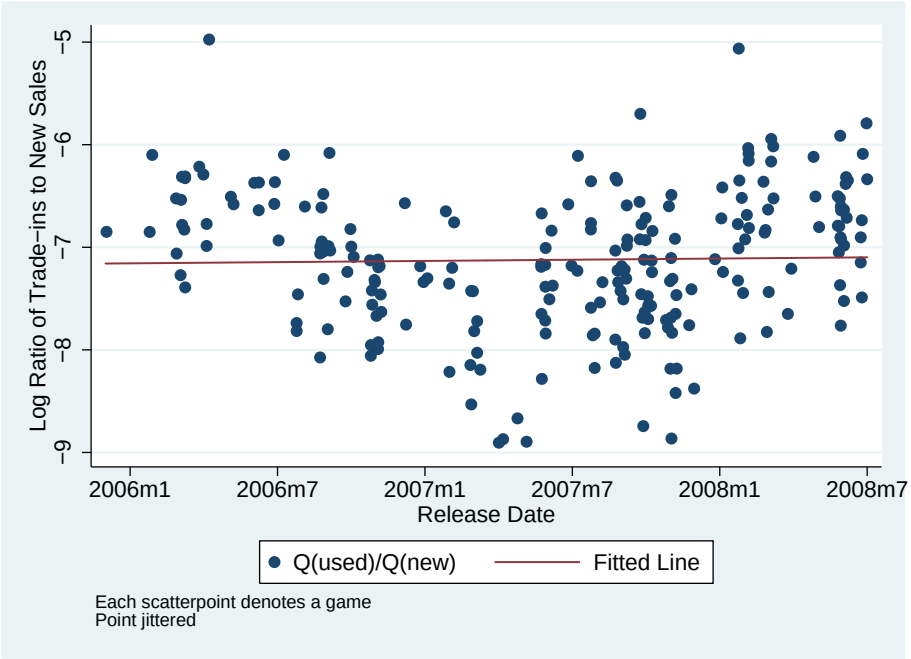

Figure 2: Ratio of Trade-ins in Raw Data to New Game Sales in Raw Data in First Two Months Following Game Release

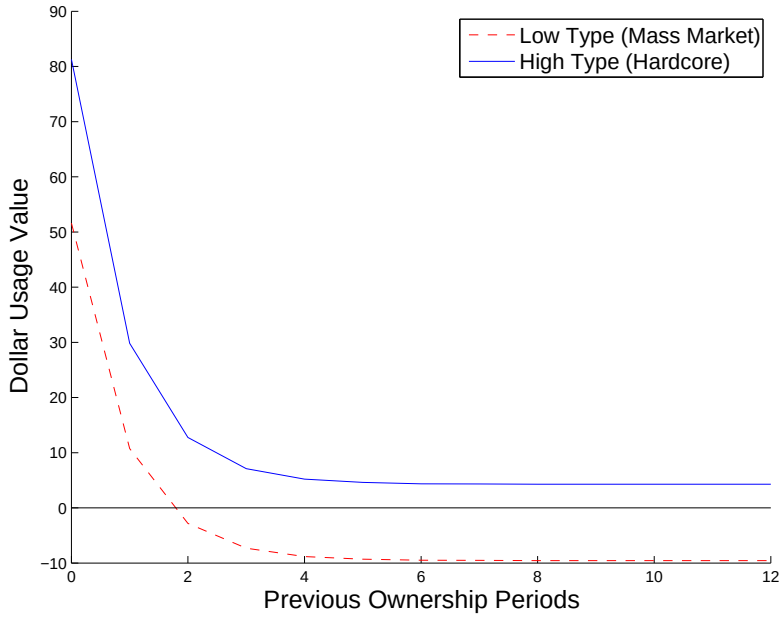

Figure 3: Static (One-Period) Dollar Value from Usage Against Length of Previous Use

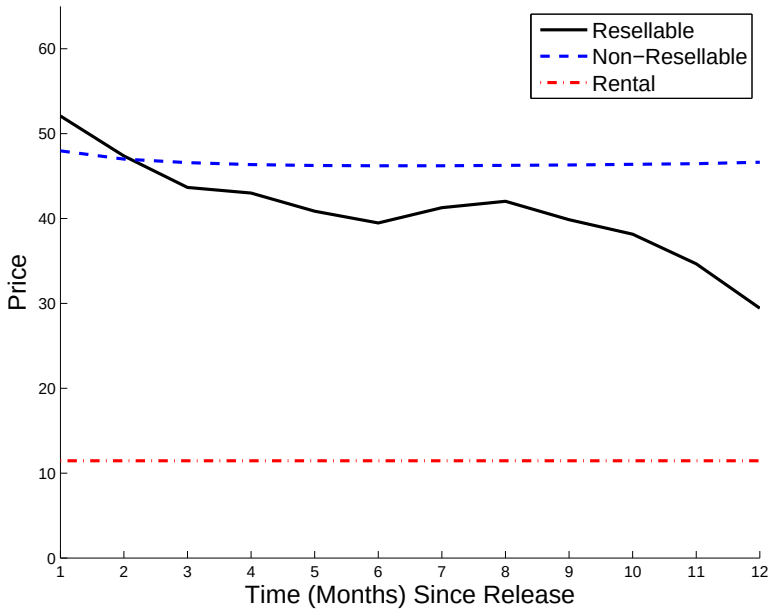

Figure 4: Simulated Price Paths Over Time Under Different Legal Environments

Table 1: Average Price and Quantity Patterns over Time

	Quantities (Normalized*)						Prices			
	New Copies		Used Trade-Ins		Used Copy Sales		Purchase Price		Trade-In Price	
Age (in Months)	Mean	SD	Mean	SD	Mean	SD	Mean	SD	Mean	SD
1	23.8	13.8	0.3	0.4	0.0	0.1	56.1	8.8	40.8	7.5
2	18.8	10.3	1.1	0.8	0.3	0.3	54.8	8.3	35.8	8.0
3	9.8	6.8	1.5	1.0	0.7	0.4	53.4	8.9	31.7	7.9
4	8.3	5.9	1.7	1.0	1.0	0.6	50.9	10.9	28.7	7.5
5	5.4	3.0	1.9	1.1	1.3	0.8	47.6	11.4	25.7	7.7
6	4.7	3.3	1.7	0.9	1.5	0.8	43.8	12.8	23.2	7.0
7	4.3	4.8	1.6	0.8	1.5	0.7	41.1	12.8	21.2	7.1
8	3.4	3.3	1.4	0.9	1.5	0.7	38.4	12.4	19.5	6.9
9	2.9	2.1	1.4	0.7	1.4	0.7	35.4	12.2	18.0	6.9
10	3.1	3.1	1.4	0.7	1.4	0.6	32.9	11.6	16.8	6.8
11	2.3	1.9	1.3	0.5	1.4	0.5	31.0	11.4	15.6	6.8
12	1.9	1.6	1.4	0.6	1.3	0.5	28.6	10.3	14.5	6.2

*Normalized by total game purchases in the first 12 months, by game.

Table 2: Price Path Regressions

	Dependent Variable is Price						
	I	II	III	IV	V	VI	VII
Lagged Price	0.962	0.960	0.955	0.955	0.954	0.950	0.945
	(0.00)**	(0.00)**	(0.00)**	(0.00)**	(0.00)**	(0.00)**	(0.00)**
Twice Lagged Price		-0.006					
		(0.77)					
2nd Critic Quintile			-0.157	-0.135	-0.185	-0.162	-0.398
			(0.62)	(0.67)	(0.58)	(0.61)	(0.24)
3nd Critic Quintile			-0.107	-0.106	-0.162	-0.106	-0.260
			(0.75)	(0.75)	(0.63)	(0.76)	(0.46)
4th Critic Quintile			0.084	0.093	0.089	0.159	-0.006
			(0.80)	(0.77)	(0.79)	(0.63)	(0.99)
Highest Critic Quintile			1.058	1.063	0.958	1.075	1.114
			(0.00)**	(0.00)**	(0.01)*	(0.00)**	(0.00)**
Month Fixed Effects				Y	Y	Y	Y
				[0.00]**	[0.00]**	[0.00]**	[0.00]**
Release Month FE					Y		
					[0.94]		
Genre Fixed Effects						Y	
						[0.58]	
Lagged Residual							0.006
							(0.79)
Constant	-0.801	-0.587	-0.625	-0.557	-0.825	-0.489	-0.031
	(0.01)**	(0.08)	(0.07)	(0.25)	(0.25)	(0.35)	(0.95)
R^2	0.91	0.90	0.91	0.91	0.91	0.91	0.90
N	2,423	2,202	2,346	2,346	2,346	2,346	2,132

P-values in parentheses or brackets (for F-statistics). * $p < 0.05$; ** $p < 0.01$

Table 3: Trade-In Price Path Regressions

	Dependent Variable is Trade-In Price						
	I	II	III	IV	V	VI	VII
Lagged Purchase Price	0.560		0.069				
	(0.00)**		(0.00)**				
Lagged Trade-In Price		0.871	0.793	0.913	0.854	0.845	0.865
		(0.00)**	(0.00)**	(0.00)**	(0.00)**	(0.00)**	(0.00)**
Lagged Trade-In Price Squared				-0.001			
				(0.02)*			
2nd Critic Quintile					0.057	0.039	-0.083
					(0.79)	(0.85)	(0.68)
3rd Critic Quintile					0.312	0.383	0.155
					(0.15)	(0.07)	(0.45)
4th Critic Quintile					0.580	0.632	0.308
					(0.01)**	(0.00)**	(0.13)
Highest Critic Quintile					1.309	1.416	0.973
					(0.00)**	(0.00)**	(0.00)**
Month Fixed Effects						Y	Y
						[0.00]**	[0.00]**
Lagged Residual							-0.079
							(0.00)**
Constant	-2.104	0.931	-0.161	0.419	0.945	0.800	0.258
	(0.00)**	(0.00)**	(0.46)	(0.14)	(0.00)**	(0.01)**	(0.41)
R^2	0.67	0.89	0.90	0.89	0.89	0.90	0.90
N	2,215	2,145	2,145	2,145	2,108	2,108	1,902

P-values in parentheses or brackets (for F-statistics). * $p < 0.05$; ** $p < 0.01$

Table 4: Estimation Results

		Value	Standard Error
Rate of Tiring Coef 1	λ_1	-1.10	0.16
Rate of Tiring Coef 2	λ_2	0.81	0.01
Hard-Core Gamers Add. Value	β	6.07	0.19
Heterogeneity Parameter 1	γ_1	0.16	0.03
Heterogeneity Parameter 2	γ_2	-0.28	0.04
Price Sensitivity Low type	$\alpha(1)$	0.37	0.05
Price Sensitivity Low type	$\alpha(2)$	0.31	0.04
Stand. Dev. (Demand Shock)	$\sigma(\xi)$	1.90	0.23
Stand. Dev. (Trans. Cost Shock)	$\sigma(\zeta)$	5.80	0.41
Corr. Btw. Dem. and Sup. Shocks	$\rho_{\xi,\eta}$	0.16	0.03

Game fixed effects, not reported above, averaged about 19.3 above ω. Their point estimates have a standard deviation of 3.2. Standard errors computed from the Hessian.

Table 5: Mean Price Elasticity Estimates

	Elasticities	
Age of Game	Static	Dynamic (3 Month)
1	−2.40	−3.13
	(0.68)	(14.20)
2	−2.53	−2.40
	(0.90)	(4.51)
3	−2.24	−1.48
	(0.67)	(0.45)
4	−2.19	−1.79
	(0.68)	(1.01)
5	−2.02	−2.04
	(0.64)	(3.20)
6	−1.91	−1.74
	(0.60)	(1.49)
7	−2.06	−1.53
	(0.65)	(0.64)
8	−2.15	−2.23
	(0.73)	(3.20)
9	−2.01	−0.90
	(0.81)	(1.06)

Standard deviations of point estimates across games in parentheses. Both types of elasticities allow consumers to update beliefs of future prices. The dynamic elasticities show the percent change in the quantity bought in the subsequent three months for a percent change in the current price.

Table 6: Counterfactual Results

Profits* (Total Across 14 Games)			
	Goods That Are:		
	Resellable†	Rented	Non-resellable
Producer Profits (in millions)	\$495	\$522	\$1,037
Q. New Sales (in millions)	12.2	83.7‡	29.3
Q. Used Sales (in millions)	2.2	N/A	N/A

*First Twelve Months Post Release. †Numbers for allowed resale are computed using observed sales. ‡Equals number times rented.

Table 7: Robustness Checks

Robustness Check Profits (In Millions)					
		Robustness Checks			
				Market Size	
	Main Results	Smaller Used Market	DF 0.95	25% Lower	50% Lower
Resellable Goods*	495	495	474	495	495
Rented Goods	522	548	373	371	213
Non-Resellable Goods	1,037	1,002	991	778	484

*Observed in Data

Table 8: Nested Logit Estimation of Substitutability

	Dependent Variable is $\log(s_{jt}/s_{0t})$					
	OLS			2SLS		
Age $(t - r_j)$	-0.194	-0.317	-0.400	-0.363	-0.624	-0.947
	(0.007)	(0.015)	(0.036)	(0.019)	(0.036)	(0.07)
Age^2		0.009	0.025		0.018	0.072
		(0.001)	(0.006)		(0.001)	(0.009)
Age^3			-0.001			-0.003
			(0.000)			(0.000)
Price	-0.012	-0.013	-0.013	-0.033	-0.034	-0.031
	(0.002)	(0.002)	(0.002)	(0.003)	(0.003)	(0.003)
Within Genre Share (σ)	0.623	0.602	0.602	-0.05	-0.068	-0.071
	(0.014)	(0.014)	(0.014)	(0.07)	(0.068)	(0.067)
Observations	2875	2875	2875	2594	2594	2594
R-squared	0.93	0.93	0.93	0.93	0.94	0.94

Notes: Standard errors in parentheses. Game and month fixed effects are included in all specifications. Regression restricted to observations where the game's age $(t - r_j) \leq 12$. Regressions of endogenous variables on instrument matrix: prices: (Rsquare = 0.92, F = 27975), within genre share: (Rsquare = 0.55, F =3220).