NET Institute*

www.NETinst.org

Working Paper #07-39

September 2007

Vertical Integration and Exclusivity in Platform and Two-Sided Markets

Robin S. Lee
Harvard Business School

Vertical Integration and Exclusivity in Platform and Two-Sided Markets

Robin S. Lee*

October 15, 2007

Abstract

This paper develops techniques to analyze the adoption decisions of both consumers and firms for competing platform intermediaries in two-sided markets, and applies the methodology to empirically measure the impact of vertical integration and exclusive contracting in the sixth-generation of the U.S. videogame industry (2000-2005). I first introduce a framework to structurally estimate consumer demand in these types of hardware-software markets which (i) simultaneously analyzes both hardware and software adoption decisions; (ii) accounts for dynamic issues including the selection of heterogenous consumers across platforms, durability of goods, and agents' timing of purchases; and (iii) explicitly provides the marginal contribution of an *individual* software title to each platform's installed base of users. Demand results show the gains obtained by a platform provider from exclusive access to certain software titles can be large, and failure to account for dynamics, consumer heterogeneity, and multiple hardware purchases significantly biases estimates. I next specify dynamic network formation game to model the adoption decision of hardware platforms by software providers. Counterfactual experiments indicate that vertical integration and exclusivity benefited the smaller entrant platforms and not the dominant incumbent, which stands contrary to the interpretation of exclusivity as primarily a means of foreclosure and entry deterrence.

Keywords: platform competition, two-sided markets, vertical integration, exclusive contracting, dynamic demand, network formation, videogame industry

JEL Classification Numbers: C61, C63, C73, L13, L14, L42, L86

*Harvard Business School and Department of Economics, Harvard University. Contact: lee54@fas.harvard.edu; 420I Baker Library, Soldiers Field Road, Boston, MA 02163. I am extremely indebted to Ariel Pakes for his advice, guidance, and support. I also would like to thank Susan Athey, Eric Budish, Marcelo Moreira, Julie Mortimer, Parag Pathak, Marc Rysman, and seminar participants at Harvard University for helpful comments and discussion during various stages of this project; furthermore, I am grateful to Marco Iansiti and Wan Wong for assistance in acquiring data, and Asi Lang, Albert Reed, and numerous others for sharing their insights into the videogame industry. Financial support from the IO Research Group at Harvard University and the NET Institute (http://www.NETinst.org) is gratefully acknowledged. All errors are my own.

1 Introduction

In most platform and two-sided markets, consumers adopt, join, or visit a platform in order to access goods or services provided by firms also affiliated with that platform. These industries include "hardware-software" markets, content and media markets, retail marketplaces, payment systems, and even some forms of "buyer-seller networks."[1] Often certain firms and their products have significant market power over consumers, inducing consumers to adopt the platform(s) that these firms have joined. Via exclusive contracts or vertical integration, platforms compete fiercely to get these firms exclusively "onboard" and dominate – if not tip – the market. This paper studies these types of exclusive vertical arrangements between platforms and firms and measures their impact on industry structure and competition.

Whether or not such arrangements are primarily pro- or anti-competitive is still a source of active debate and an open empirical question. On one hand, exclusive contracts raise antitrust concerns since they may reduce competition by deterring entry or foreclosing rivals.[2] The presence of network externalities prevalent in these types of markets heightens this fear: not only does exclusive access to the services or goods of a particular firm draw more consumers onboard a platform, but in a dynamic and networked context this effect may be amplified.[3] Furthermore, from a consumer welfare perspective, these vertical arrangements may limit consumer choice by preventing consumers on competing platforms from accessing exclusive content, products, or services.

On the other hand, exclusive arrangements are also argued to have pro-competitive benefits. Standard theoretical justifications include the encouragement of investment and effort provision by contracting partners.[4] Other possibilities also emerge within networked environments. In nascent markets, integration into one side of the market by a platform provider may be effective in solving the "chicken-and-egg" coordination problem, one of the fundamental barriers to entry discussed in the two-sided market literature.[5] Furthermore, exclusivity may be an integral tool used by

[1] Examples of hardware-software markets include CD and DVD players and their respective media, videogame consoles and games, and PCs and their compatible applications or peripherals (where hardware usually refers to the CPU and OS combination). Many standard battles occur within these industries, whereby the success of competing hardware platforms (e.g., VHS/Betamax, HD-DVD/Blu-ray) depends crucially on the provision of compatible software. Other examples include: cable/satellite television, satellite/terrestrial radio, internet portals/aggregators, and online music services; shopping malls, department stores and brick-and-mortar merchants who "rent" space to vendors and manufacturers; credit and debit cards (Visa, Mastercard) and electronic money systems (Paypal, FeliCa, Edy); and HMOs and hospitals as buyer-seller networks, whereby consumers join the former to access the services of the latter. Online advertising exchanges (Google, Microsoft, Yahoo!) are also examples, whereby the relevant parties are advertisers who join to get access to "publishers," or sites where ads are shown.

[2] See, e.g., Mathewson and Winter (1987), Rasmusen, Ramseyer, and Wiley (1991), and Bernheim and Whinston (1998). Whinston (2006) and Rey and Tirole (2007) provide excellent overviews of the theoretical literature on vertical foreclosure.

[3] E.g., exclusivity induces more consumers to join a platform, which in turn encourages more firms to subsequently join, and then more consumers, and so on. Thus, an entrant platform may not be able to gain critical mass and hence be deterred from entering (Shapiro, 1999). Evans (2003) discusses other related antitrust issues brought up within two-sided markets. Accessible surveys of the network effects literature in general are Katz and Shapiro (1994), Shy (2001) and Farrell and Klemperer (2007). Empirical papers on measuring the impact of indirect network effects include Saloner and Shepard (1995), Gandal, Kende, and Rob (2000), and Rysman (2004).

[4] See, e.g., Marvel (1982), Klein (1988), Besanko and Perry (1993), and Segal and Whinston (2000).

[5] I.e., without firms, consumers will not join a platform; however, without consumers, firms will not join. See e.g.,

entrant platforms to break into established markets: by preventing contracting partners from also supporting the incumbent, an entrant may spur adoption of its own platform and thereby spark greater platform competition.[6]

Given the growing prevalence of networked and platform industries, the ability to resolve this theoretical ambiguity is of central importance for policy and regulation. Not only were the competitive implications of integration and exclusive contracting at the heart of recent prominent antitrust cases in networked industries – *U.S. v. Microsoft* [253 F.3d 34 (2001)],[7], *European Union v. Microsoft* [COMP/C-3/37.792 (2004)], and *U.S. v. Visa* [344 F.3d 229 (2003)][8] – but they are also the central issues to consider whenever evaluating the impact of exclusive carriage deals in the media industry[9] or determining whether or not to open up closed hardware-software systems to competitors.[10]

This paper has two primary goals. The first is to provide a framework for analyzing the adoption decision of both consumers and firms for competing platform intermediaries. Fundamentally, any analysis of how exclusive arrangements affect a platform market's competitive structure requires an understanding of how parties on each side of the market choose which platform to join. The challenge is to account for the inherent dynamic nature of a networked industry, whereby each agent makes a choice anticipating the future actions of others. A consumer times her purchase and chooses a particular platform anticipating the adoption decisions of other firms and consumers (as well as the future paths of prices and qualities); similarly, a firm joins a platform only after forming expectations over the number of consumers and other firms that will also come onboard.

The second objective is to apply the methodology to empirically measure the impact of exclusive contracting and vertical integration by hardware providers into software provision in the sixth generation of the US videogame industry (2000-2005). As a canonical hardware-software market comprising three differentiated hardware platforms each with its own distinct base of software, the videogame industry exhibits features easily generalizable to a variety of networked environments – indeed, two of the main platform providers, Sony and Microsoft, are themselves participants in myriad other platform environments. The videogame industry is also convenient to study since

Caillaud and Jullien (2003), Armstrong (2006), Rochet and Tirole (2003, 2006).

[6]Lee (2006) provides a related model of platform competition whereby a large enough marginal contribution of a single firm is sufficient to prevent the market from tipping completely to one platform.

[7]See Whinston (2001) for a discussion on the Microsoft case, which involved Microsoft's integration of Internet Explorer into Windows and the possible foreclosure of Netscape, a rival browser application provider.

[8]The Visa case dealt with the possible foreclosure of competing payment cards (American Express and Discover) from accessing the services of exclusive member banks; Rochet and Tirole (2002) provide an economic analysis and discuss the potential antitrust concerns within the payment card industry.

[9]E.g., in 2007 Major League Baseball and satellite television operator DirecTV agreed to an exclusive deal which would deny cable television subscribers access to a package of out-of-market games. The deal was eventually scuttled after a U.S. Senate hearing into the issue pressured MLB to renegotiate with excluded parties; MLB ultimately agreed to supply television feeds of out-of-market games to both cable and satellite TV operators under threat of losing its anti-trust exemption ("Baseball Strikes Deal to Keep 'Extra Innings' Package on Cable," *Associated Press*, April 4, 2007).

[10]E.g., France recently passed a law that allowed regulators to require Apple to open up its iTunes music service to other companies' music players ("French iTunes Law Goes Into Effect," *Associated Press*, August 8, 2006). Consumer rights organizations in other countries including Germany, Finland, and Norway as well as the European Union consumer chief have taken similar stances ("EU takes aim at Apple over iTunes," *Reuters*, March 11, 2007).

exclusionary contracting or integration by platforms into software development is not intended to foreclose other third-party software providers; on the contrary, any exclusive titles are intended to attract other software developers as much as they are intended to attract consumers.[11] As a result, focusing on the videogame industry can separate out the possible foreclosure effects in software provision and instead focus on foreclosure at the platform level.

The main finding of this paper is that exclusive arrangements and vertical integration benefited the two smaller entrant platforms at the expense of the industry incumbent, and were generally pro-competitive at the platform level.[12] Absent the ability to produce or acquire exclusive titles, competitive conditions would have been vastly different: the incumbent would have captured an even more commanding share of the market than it did, with the two entrants increasingly marginalized. As such, the ability to integrate and engage in these exclusive arrangements may have prevented monopolization by a single hardware provider.

However, when exclusive vertical arrangements are prohibited, consumers are shown to possibly benefit from access to a greater selection of software titles onboard any given platform *holding fixed* the pricing and entry/exit decisions of all platform providers and software titles – consumer welfare would be predicted to increase by about $7B during the five year period analyzed. However, this is partially a static result; whether or not consumers would have truly been better or worse off depends crucially on how software production would have responded to the lack of integration or exclusive dealing, and how the dominant platform would have reacted to its increased market power. If either previously integrated "first-party" titles no longer are produced, or if monopolization induced the exit of competing platforms (even without raising prices), any potential consumer welfare gains from increased access to software would be mostly eliminated.

Even though the impact on consumer and hence total welfare is ambiguous, the analysis nonetheless robustly predicts that lack of these exclusionary vertical agreements would concentrate market power and yield one dominant platform provider. This highlights the importance of accounting for dynamic consequences (increased platform competition and investment) when addressing the potential static costs and inefficiencies (reduced access and software variety onboard each platform) of exclusive arrangements.

When considering this implication in a broader context, it is important to note that "forced exclusivity" contracts – whereby a software developer would not be allowed to release software for a hardware platform unless it did so exclusively – are implicitly not considered when evaluating the effects of exclusivity. During most of the 1980's and 1990's, the dominant videogame platform provider, Nintendo, used to force its developers to release games exclusively via these contracts. However, following a 1992 antitrust investigation related to *Atari Games Corp v. Nintendo of*

[11] Unlike the computer application space whereby users typically only need one word processor, browser, or media player, videogames are not direct substitutes for one another and are instead more like "disposable" media goods which face demand for continual replacement. Thus exclusive software, by drawing onboard more consumers, increases the potential market that another software providers can access but does not crowd out or compete against these other titles.

[12] One of the "entrants" (Nintendo) was itself a veteran of the industry, but entered (alongside Microsoft) with its sixth-generation console a year after the dominant platform provider (Sony).

America, Inc [975 F.2d 832 (1992)], Nintendo dropped these practices.[13] Since then these contracts have not been utilized or observed within the videogame industry, and will not be considered within the space of vertical agreements considered here.[14]

Previous empirical work on measuring the effects of exclusive dealing and foreclosure has primarily focused on supply-side efficiencies and the possibility of "upstream" foreclosure. Asker (2004) and Sass (2005) analyze the relationship between manufacturers and distributors in the beer industry, and find evidence that exclusive arrangements seem to enhance efficiency, improve distributor effort, and do not foreclose smaller brewers. Chipty (2001) studies the integration between programming and distribution in the cable television industry, and finds that although certain programming may have been foreclosed from distribution by integrated cable providers, the associated efficiency gains have likely offset any social costs generated.

In contrast, this paper focuses on the competition between "downstream" platforms, and primarily how integration and exclusivity interacts with the networked aspect of the industry in attracting both software and consumers to join. I examine a counterfactual regime whereby platforms are unable to integrate into software development, and exclusive contracts are prohibited. To do so, I determine both demand (consumer) and supply (software) responses to such changes. This paper is the first to my knowledge to estimate dynamic demand for multiple "sides" in a platform market, and explicitly account for the rematching process between contracting partners within a counterfactual regime.

1.1 Framework for Analysis

The paper is structured in three parts: the first two develop and apply a framework for analyzing how industry participants chose which hardware platform to purchase or develop for, and the third applies estimates obtained from these first two stages to evaluate regimes in which exclusionary vertical contracts are prohibited.

Consumer Demand

The first part of this paper focuses on analyzing consumer demand, and develops a methodology for structurally estimating a demand system in general platform-intermediated markets. Although also of independent interest, the primary purpose of this section is to determine how consumer demand for hardware responds to changes in software availability, and what consumer demand for any given piece of software would be conditional on that title developing for any set of platforms. These estimates are used as inputs in the second part of the paper, which focuses on determining which platform(s) a software title will join and hence requires an estimate of each title's expected profits; this in turn requires a prediction of not only how many people onboard a platform will

[13] Shapiro (1999), Kent (2001).

[14] These types of forced agreements have also been ruled to be anti-competitive in other industries: e.g., the courts in *U.S. v. Visa* found these types of contracts to "weaken competition and harm consumers" by limiting output of rival payment card providers and foreclosing them from competing in related markets.

purchase a particular title, but also how many more consumers will purchase the *platform* itself.

In networked industries, a consumer's choice of platform in is a function of not only the platform's own characteristics, but primarily over the goods or services that are or will be available onboard the platform. However, different consumers will choose different hardware platforms based on their preferences over affiliated products: just as consumers choose a local community to best satisfy their preferences as in Tiebout (1956)'s model of local expenditures, so do they behave with respect to selecting a particular platform or hardware device.[15] Consumers who have purchased a platform are more pre-disposed to purchase those products onboard; failing to account for this selection will lead to significant upward biases in estimates of the quality and contribution of goods and services to that platform.

In response to this challenge, I explicitly account for this form of interaction between the hardware and software sides of the market and develop a structural model of consumer demand which: (i) simultaneously estimates both (hardware and software) sides of the market; and (ii) accounts for dynamic issues including the selection of heterogenous consumers across platforms, durability of goods, and agents' timing of purchases. I leverage the dynamic aspect of my panel data and use the differential responses of unaffiliated and affiliated consumers to changes in software availability over time in order to identify the selection of consumer heterogeneity across platforms.

Via the demand system, I am also able to recover the marginal contribution of an *individual* software title to each platform's installed base of users. This point is one of the primary reasons for adopting a structural approach in demand estimation. In these types of platform markets when one side is oligopolistic and the identity of contacting partners matters, analyzing the impact of exclusive arrangements requires computing how profits are affected when *individual* contracting partners change.[16] Although other papers have estimated demand systems in platform markets, none have done so without ignoring or taking a reduced form approach to one side of the market, and consequently have been unable to address both this concern as well as the selection of consumers across platforms.[17]

Additionally, accounting for dynamic concerns is critical in platform markets as the affiliation decision often involves the purchase of a durable good or agreement to a long-term commitment. A large literature has shown the limitations of applying a static methodology to a dynamic setting.[18]

[15] Dubin and McFadden (1984) study a similar issue, in which household demand for electricity depends on the choice of space or water heating.

[16] See e.g. Lee (2006). This is similar to calculating the underlying "value" functions in multilateral contracting environments with externalities (e.g., Prat and Rustichini (2003)), partition functions in coalitional bargaining games, and graph functions in bipartite network formation games (e.g., Bloch and Jackson (2007)).

[17] E.g., Song and Chintagunta (2003), Clements and Ohashi (2005), Prieger and Hu (2006), and Corts and Lederman (2007) assume software utility on a platform is simply a function of the total number of titles on a system; Nair (2007), although accounting for dynamic concerns, analyzes demand only for software and does not consider the hardware side.

[18] E.g., a consumer does not purchase a durable good that she has already purchased in the past, and upon purchase does not (usually) keep participating in the market; a consumer may purchase items in order to stockpile goods; and a consumer may have expectations over future product availability and quality due to technological improvements, and delay purchase not because current goods are inadequate, but because future goods may be more preferable. See e.g., Melnikov (2001), Song and Chintagunta (2003), Hendel and Nevo (2006), Carranza (2006), and Gowrisankaran and Rysman (2007) for other applications of modelling dynamic demand.

To account for these dynamics in hardware and software adoption, I adapt and extend a number previously introduced techniques – including those pioneered in Rust (1987), Berry (1994), and Berry, Levinsohn, and Pakes (1995), and later synthesized in a dynamic context in Hendel and Nevo (2006), Gowrisankaran and Rysman (2007) and Nair (2007) – in order to formulate and solve each consumers' dynamic optimization problem in a platform environment. The extensions include the explicit handling of seasonality effects (crucial for most consumer product industries), the persistence of unobservable product characteristics, and the use of a more general evolution process for product qualities and prices.

Using a new dataset containing monthly level aggregate sales and prices for every hardware platform and software title in the sixth generation of the videogame industry, I apply the methodology to estimate consumer demand. Results show that the gains obtained by a platform provider from exclusive access to certain software titles can be large: although most titles do not have any impact on hardware sales, certain hit titles may be able to increase the installed base by 10% or more. In general, failure to account for dynamics, heterogeneity, and multiple purchases by consumers significantly biases the predicted impact a title has on the demand response of consumers.

Hardware-Software Network Formation

The second half of the paper examines the contracting decisions of software firms and hardware platforms. In analysis of consumer demand, the set of software products for each platform and at each point in time has been conditioned on in the data. However, when institutional features of an industry change – as is the case when hardware platforms cannot vertically integrate into software provision or offer contracts contingent on exclusivity – it is unlikely that the contracting relationships between parties will remain the same. To determine the supply-side response of which platforms each software title joins, I define and compute a new equilibrium for a dynamic network formation game whereby software titles are allowed to freely choose which platforms to develop for, having observed the previous actions of consumers and other industry agents and having formed expectations over the future profitability of each potential strategy.[19] The setup allows for contracting partners and consumer demand to change over time with past actions influencing future decisions, a key feature of any networked industry. Although each software title makes only a single decision, each employs an equilibrium strategy depending only on the value and expected evolution of certain payoff relevant state variables; furthermore, in equilibrium these expectations over the variables are assumed to be consistent with their actual realizations. In this regards, this model is similar in spirit to Ericson and Pakes (1995) and the literature that follows.

For the equilibrium computation, I require an estimate of the profits each software title expects to receive if it chose to develop for any set of platforms. As discussed earlier, the expected number of consumers who will purchase a given title is obtained from the first half demand estimates. However,

[19] Here, when platforms are forbidden to engage in exclusive deals, they are treated as passive agents. See Hagiu and Lee (2007) for a model of platform competition whereby platforms actively bargain for (exclusive) content in media markets.

construction of profits also requires knowledge of the underlying "porting" costs of developing for different sets of platforms. These are unobserved. To address this issue, I first assume each software title chose the set of platforms which maximized its expected profit, and compute the difference in profits (as a function of costs) each title would have expected to receive if it had instead chosen to develop for a different set of platforms. I then recover an estimate of porting costs that rationalize each title's observed choice via an inequalities-based estimator developed in Pakes, Porter, Ho, and Ishii (2006). To evaluate the fit of the dynamic network formation model and estimated parameters, I initially fix the decisions of exclusive first-party titles and compute a new equilibrium among remaining third-party titles. I find that the specified model accurately predicts platform installed base figures, market shares, and contracting decisions for "hit" titles.

Counterfactual Experiments

Finally, I analyze a counterfactual regime in the absence of vertical integration and exclusive contracting between hardware platforms and software titles, and compute the new predicted equilibrium outcomes. To account for the possibility that eliminating exclusive arrangements may have affected the production of first-party titles, I compute two separate counterfactual regimes whereby first-party titles are either assumed to enter as third-party products, or are eliminated altogether. In both cases, counterfactual simulations indicate that the industry is far more competitive when exclusive vertical arrangements are allowed than when they are prohibited.

1.2 Road Map

In the next section I describe the U.S. videogame industry, the role of vertical integration and exclusivity, and important features and stylized facts of the market which must be captured for any reasonable analysis. Section 3 overviews the theoretical issues involved with demand estimation in platform markets, and develops the full dynamic model with the accompanying details on estimation, inference, computation, and identification; estimation results are presented in section 4. I lay the groundwork for computing the dynamic network formation game in section 5, and also discuss how to recover the underlying porting costs borne by software firms. Finally, I analyze counterfactual regimes whereby exclusive agreements are banned in section 6, and conclude in section 7.

2 Application: The U.S. Videogame Industry

2.1 Industry Description

Starting as a fringe industry in the early 1970's with the introduction of a home version of *Pong*, the U.S. videogame industry has since grown to reach $13.5B in revenues in 2006.[20] Increasingly,

[20] Approximately $6.5B from console videogame software, $1B from PC videogame software, and the rest from console hardware and accessories. *Entertainment Software Association 2006 Sales, Demographic and Usage Data.*

as evidenced by the widespread adoption of the new generation of consoles introduced in 2006, videogames have broadened their appeal and user base from a child's hobby to something more mainstream: 69% of American heads of households engage in computer and videogames with the average age of a player being 33 years old,[21] and market penetration of videogame consoles reached 41% of US television households (45M) in 2006.[22]

A videogame system comprises a hardware console (the platform) and software (its games). In the current and most recent generations, each console is and has been provided by one firm – the platform provider – as a tightly integrated and standardized device which is required in order to run any of the titles provided for the system.[23] Videogames, on the other hand, are brought to market by two types of entities: *developers*, who undertake the programming and creative execution of each title; and *publishers*, who handle the marketing and distribution of games. Publishers are usually integrated into software development and have their own in-house development studios; although independent software development studios do exist, as the costs of developing games have increased over time – with average costs reaching $6M during the late 1990's – these studios still must often turn to software publishers for financing in exchange for distribution and publishing rights.[24]

Console manufacturers have also integrated into software publishing and development, with each platform provider typically having its own publishing unit and game development studios. Any title produced by the console maker via its own integrated studios or acquired and distributed by its own publisher is known as a *first-party* title, and is exclusive to that hardware platform. All other games are *third-party* titles and are produced by other firms. Within a generation, games developed for one console are not compatible with other consoles; in order to be played on the other console, the game must explicitly be "ported" by a software developer and another version of the game produced.[25] Consequently, the choice of which platforms to develop for is highly strategic: a third-party software developer can release a title on multiple platforms in order to reach a larger audience and pay additional porting costs (which may include development and distribution costs) or develop exclusively for one console and forego consumers on other platforms. Even in the latter case, the developer has multiple options: it can enter into a publishing agreement with the console provider, or opt to sell its game or even entire studio outright.

Since videogame consoles usually have little if any stand alone value, consumers purchase a particular console only if there are desirable software titles on that system.[26] At the same time, software publishers release titles for consoles that either have or are expected to have a large installed base of users who potentially will purchase their games.[27] These cross-side network externalities

[21] *2006 Essential Facts*, Entertainment Software Association.

[22] *The State of the Console*, Nielsen Media Research. March 5, 2007

[23] NB: this is unlike the PC industry where a system's hardware and operating system may be provided by different firms, and can be modified and configured further by the end user.

[24] Coughlan (2001).

[25] A notable exception is "backwards compatibility," whereby a new console can play the games made for the previous generation of that particular console. For example, each of Sony's three home videogame consoles could play games made for the previous generation.

[26] Recently however, consoles have been able to perform more tasks such as watch DVDs, access the internet, and purchase digital content or services.

[27] As primarily a fixed costs business, maximizing the potential audience for a title is one of the primary ways to

and "two-sidedness" are manifest in most hardware-software industries, and in particular yields a complex form of competition between rival platforms. In particular, on the pricing front most videogame console manufacturers subsidize the sale of their hardware devices to consumers, selling them close to or even below cost, and make profits by charging publishers and developers a royalty for every game sold on that platform.[28]

As discussed in the introduction, as the dominant videogame platform provider during most of the 1980's and 1990's, Nintendo used to write forced exclusivity contracts with developers, committing them to 2-year exclusive deals in exchange for the right to develop for its system. This changed following a 1992 antitrust investigation related to *Atari Games Corp v. Nintendo of America, Inc*, and Nintendo dropped these practices. Since then, such forced exclusivity contracts have not been observed. In their place, console manufacturers have primarily relied on internal development, integration (i.e., outright purchase), or favorable contracting terms to third party developers or publishers (e.g., lower royalty rates, lump sum payments, or marketing partnerships) in order to secure exclusive titles. More recently, as games have been becoming more and more expensive to develop, most third-party titles have chosen to multihome in order to access a wider audience to recoup these fixed costs; consequently, console providers have become more reliant on their own internal first-party titles to differentiate their platforms.[29]

2.2 The Sixth Generation: 2000 - 2005

The videogame industry typically witnesses the release of a new set of consoles approximately every five years. Since hardware capabilities remain fixed within a generation to ensure compatibility and standardization, it is only during these generational shifts that new hardware with more powerful processing power and graphical abilities can be introduced. In October, 2000, Sony released its Playstation 2 (PS2) console, among the first of what has since been referred to as "sixth-generation" of videogame consoles.[30] The PS2 was a followup to the original Playstation (PS1), Sony's wildly successful entry in the previous generation.[31] Sony also had the advantage of being the first out of the gate with a sixth-generation hardware; only a year later did industry veteran Nintendo release its Gamecube (GC) and new entrant Microsoft bring its Xbox console to market. By the time the first seventh generation console was introduced in October 2005, Sony's new console would have gone on to sell almost double the number of hardware devices as both its competitors combined.

This paper's focus on the sixth-generation is for several reasons. First, it marked the arrival of a new competitor – Microsoft – to the industry. Itself a veteran and competitor in other myriad

ensure development costs are recouped.

[28]See Hagiu (2006) and Evans, Hagiu, and Schmalensee (2006) for more on this point.

[29]Instead of complete exclusivity for third-party titles, there has recently been a greater use of "timed" exclusives or exclusive "premium" add-ons to differentiate a title across platforms – the former allowing the console to have a particular title exclusively for a set period of time (usually 6 months), the latter only becoming feasible with the rise of downloadable content.

[30]Sega's Dreamcast, released a year earlier, was discontinued on January 31, 2001 and is not considered in this paper.

[31]The PS1 had sold over 70M units by that point, and would go on to sell over 100M units by the end of 2005. (`http://www.playstation.com/business.html`)

platform industries, Microsoft's entry tactics included the acquisition of several software developers; whether or not Microsoft would have been able to gain a foothold into the industry absent integration is an open question. Secondly, the three platform providers are the same as the current (seventh) generation providers, providing timeliness to this line of inquiry. Finally, this particular generation marked the first steps towards placing the videogame industry firmly within the convergence battle between personal computers and other general consumer electronics. Starting with the introduction of DVD and online capabilities with these consoles, videogame platform providers have added significant non-gaming functionality to their devices such that they now function as fully independent media hubs and thereby placed their consoles as platforms in several other markets.[32] Consequently, these success or failure of these hardware devices vis a vis one another has a dramatic impact on industries far removed from just videogames.

2.3 Data and Descriptive Statistics

The main data comes from a new data set, and consists of monthly observations from January 1994 to October 2005. Each observation includes the average selling price and quantity sold of each videogame console, the average selling price, quantity sold, and other descriptive information for each software title on each console (including ratings provided by the industry *Entertainment Software Rating Board*, genre, and date of release). Price, quantity, and software descriptive data for videogame titles and consoles was obtained from the NPD Group, a market research firm, for the dates specified.[33] Prices are adjusted via the Consumer Price Index.[34] In the data, there is information on a total of 12 consoles, two of which existed prior to 1994 and the other 10 having been introduced during the period of the data, and 6606 videogame titles. However, since the analysis focuses on the sixth-generation, the data utilized is selected from the period between September 2000 and October 2005. During this period, three videogame consoles and 1581 unique software titles were released. The population of potential consumers is given by the number of television households collected on a yearly basis from Nielsen and linearly interpolated to the monthly level.

General descriptive statistics for each of the three sixth-generation consoles are provided in table 1. What follows are additional stylized facts about the industry:

- *Prices:* Hardware prices followed very steady paths, interrupted only by two major discrete jumps. The PS2 and Xbox started retailing for \$300, but in May 2002 both simultaneously cut their prices by \$100 prior to the "E3" industry trade show. Nintendo followed with a \$50 price cut of its own from \$200 to \$150. Microsoft and Sony again quickly dropped prices one after the other two years later. Figure 1 illustrates the price paths of each hardware console. Software prices, however, follow much more regular price drops, with price cuts usually following the first few months of a new title's release.

[32] E.g., these new consoles are also able to play next-gen HD-DVD/BluRay discs, download movies and digital content, watch IPTV, and network with other users across the internet.

[33] The information is collected from approximately two dozen of the largest retailers in the US, which account for about 65% of videogame and console sales, and is extrapolated by NPD for the entire US market.

[34] All urban consumers, all items less food and energy.

- *Seasonality:* The videogame industry, like most consumer product markets, exhibits considerable seasonality both in consumer demand as well as in software supply. Figure 1 also shows the number of total hardware consoles sold each month, and during holiday months (November and December) the number of consoles sold is easily double or triple the average number sold in other months. Software supply also exhibits significant variance across months: some have over 100 new titles released across systems, and others less than 5.

- *Exclusivity and Multihoming:* There is significant variation in the number of software titles that are exclusive across platforms: although nearly 64% of all unique software titles are exclusive to one console, the majority are located on the PS2. On the other hand, the majority of GC tiles are available on the other two console systems. At the same time, as seen in Table 2, the top titles for the GC are all exclusive whereas the PS2 has more nonexclusive titles in its top sellers.

- *Concentrated Software Sales:* Videogames, like motion pictures, are primarily a hit driven industry whereby most sales are concentrated among a few top-selling games. Despite there being over 1500 unique titles released over the three sixth-generation consoles, the top 10 listed in Table 2 on the PS2, Xbox, and GC accounted respectively for 18%, 33%, and 47% of total platform software sales. The concentration applies to software publishers as well – with over 150 publishers who have released a game for a sixth generation console, the top 5 were responsible for 50% of all software sold (by total quantity sales), and the top 20 responsible for over 90%. Indeed, only 5 publishers command greater than 10% market share on any given console, and 3 of them are the console manufacturers themselves: Sony, Microsoft, and Nintendo. Finally, the concentration of software sales occurs early in a title's life, with on average over 50% of total sales occurring within the first 3 months of release.

- *Significant Consumer Heterogeneity:* According to Nielsen, the heaviest 20% of videogame players account for nearly 75% of total videogame console usage (by hours played), averaging 345 minutes per day during the fourth quarter of 2006. At the same time, the fastest growing segment of users are known as "casual gamers" who spend less than 5 hours a week playing games.

3 Consumer Demand

In order to study the impact of a policy restriction on exclusivity, an understanding of how consumer demand responds to changes in software availability is required. Not only is this used to compute the resultant platform and software market shares following a change in industry structure, but it also is needed to understand the incentives governing each software title's decision of which platform(s) to develop for in the first place. This section develops a structural model of consumer demand for both hardware and software in general platform markets, and then applies it to the videogame industry.

I begin in 3.1 by overviewing the theoretical issues involved with consumer demand estimation in platform markets, and introduce the main structural innovations of the paper within a static environment. These include endogenizing platform utility as a function of affiliated software products, and handling the selection of heterogeneous consumers across platforms. I discuss why standard discrete choice methods are not sufficient to capture realistic substitution patterns or elasticities of demand within these markets, why they are not able to recover the quality or impact on demand of an individual software title.

In 3.2, I present the full dynamic model used in estimation which not only allows for heterogeneous consumers and multiple hardware system purchases, but also for consumers to be forward-looking agents: these consumers account for future software releases and pricing or quality changes, and may delay purchase in anticipation of receiving higher utility in the future. The major assumption I use to model consumer expectations is similar to but more general than the one used in Melnikov (2001), Hendel and Nevo (2006), and Gowrisankaran and Rysman (2007): rather than explicitly model supply-side decisions (which includes software availability and pricing) when evaluating consumer expectations, I instead assume consumers use a reduced-form approximation of the evolution of each product's quality which depends on past values of itself and of other competing products. These expectations are assumed to be rational in the sense that they are consistent with the actual distribution of realized product qualities, and may be seen as reasonable approximations of consumer beliefs.

I also make the additional assumption that software titles are imperfect substitutes, and each title effectively competes within its own market. This assumption is not unreasonable for this particular industry: for example, Nair (2007) finds empirically that videogames are generally not substitutable for one another in his analysis of videogame software.[35] Considering there are a large number of titles (even within a particular genre), each with their own distinct idiosyncracies, plots, characters and style of play, such reduced form results indicating no first-order substitution effects across titles are not particularly surprising. The important implication this particular assumption delivers is that each consumer, after solving her appropriate dynamic policy for hardware purchase, can solve an independent optimal stopping problem for each individual piece of software. Furthermore, even if there exist some minor cross-title effects between certain titles, they are unlikely to affect demand estimates significantly or change any of the results in the counterfactual exercises conducted later in this paper.[36]

In 3.3, I develop the estimation and associated computational routine. Those not interested in details may skip ahead to 3.4 where identification of the model is discussed, or to section 4 where

[35]For software released between 1998-2000 on Sony's original Playstation console, Nair (2007) shows cross-price effects across games to be low (even when accounting for strategic release timing on the part of game developers), consumers do not seem to exhibit intertemporal substitution within genres, entry by hit games do not have a significant effect on sales or prices of games within a genre, and rates at which game prices fall are independent of competitive conditions within the market.

[36]As will be evident, only the contracting decisions for "hit" titles (those of sufficiently high quality) significantly affect platform market shares. Such titles have significant market power and are the least likely to be affected by any potential competition from other titles.

results from the demand estimation are presented.

3.1 Demand Estimation in Platform Markets

Assume there are J hardware platforms and K software products available in a given market, where J and K will refer to both the set and number of each respective product. Let $\mathbf{K}_j$ represent the set of software available on platform j. A consumer can only utilize a software product $k \in \mathbf{K}_j$ if she first purchases platform j.[37] For now, assume consumers may only purchase one hardware platform and the environment exists only for one period in order to abstract away from dynamic concerns. The timing of actions is as follows:

> Stage I: Consumers may choose to purchase any hardware platform $j \in J$.
>
> Stage II: If consumer i has chosen platform j, she may purchase any subset of products $K_{i,j} \subset \mathbf{K}_j$

I focus on the problem of estimating a demand system for consumer behavior observed in stages I and II. I assume the econometrician observes the aggregate share of each hardware platform chosen in Stage I and share of consumers on each platform who buy each piece of software in Stage II. In this platform setting, there are two primary issues that a demand system must capture:

i. *Platform utility is endogenous and a function of affiliated products* $\mathbf{K}_j$: a consumer derives utility from purchasing a particular piece of software, and this must be accounted for in the utility she expects to derive from adoption of the platform. Any parameters that enters into the specification of utility of both software and hardware – e.g., price sensitivity – should be consistent and jointly estimated. Furthermore, a consumer's utility upon joining a platform can only be a function only of those products affiliated with that platform, and the choice set over software products changes depending on which platform she joins.

ii. *Consumers select across platforms according to their preferences and characteristics*: failing to account for either heterogeneity in consumer preferences or their selection across platforms will lead to biased estimates of the quality and contribution of a piece of software to consumer utility, since those onboard the platform have already exhibited their preference for those goods affiliated. Consequently, any model which implies consumers who purchase and those who do not are identical is likely misspecified. Nonetheless, this is often the assumption made when estimating software demand without also explicitly accounting for hardware demand.[38]

[37] Each software title k may be available on multiple platforms – i.e., $\mathbf{K}_j \cap \mathbf{K}_{j'}, j \neq j'$ need not be empty. As long as the econometrician can distinguish from which platform a purchaser of k owns – which is the case in this paper – the following discussion does not change.

[38] To preview later discussion, the second issue is amplified when demand is dynamic since there is an additional dimensionality of selection over time. In many hardware-software industries, early adopters tend to be those who exhibit high values for software or low price sensitivity. Without heterogeneity, the characteristics of outside consumers (non-purchasers) do not change over time, and any model would have difficulty rationalizing improving product quality and declining prices with falling shares of observed purchasers.

I have not yet explicitly specified when $\mathbf{K}_j$ is determined. It may be the case that all software products join a platform prior Stage I, all join immediately before Stage II but following Stage I, or some join before Stage I and others before Stage II. In the case whereby consumers perfectly observe the set of software products available on each platform $\mathbf{K}_j$ and perfectly know their utility over each software product prior to platform adoption in Stage I, then there is no need to separate out the consumer's decision into these two stages: consumers essentially choose "bundles" of both hardware and software simultaneously, and thus any discrete choice demand framework over properly specified bundles of goods would be adequate.[39] Yet, in many environments consumers do not have complete information about either the identity of software products available, or the utility they derive from them.[40] Insofar there is consumer uncertainty *about software quality or availability* prior to Stage I which is only resolved after the platform has been chosen, then there is a need to link software demand following the realization of uncertainty with the ex ante expected utility from software in hardware demand. For now, I assume $\mathbf{K}_j$ is exogenously given.[41]

Previous work on estimating demand in platform markets have usually taken one of two approaches: (i) estimate only one side of the market (typically the hardware side) using a reduced form approximation for the contribution of utility of the other side (which usually is the number of complementary products available);[42] (ii) estimate each side in a separate two-stage procedure, combining the software estimates in the first stage to construct a measure for hardware utility. Limitations of the first approach are numerous, notably its inability to estimate the entire structure of demand for a given market and to conduct counterfactual experiments in which contracting partners change. Furthermore, a "killer-application" developed for a hardware system or world-class cancer treatment center as part of an HMO group are not replaceable by a handful of mediocre titles or mid-tier hospitals; yet this is the equivalence suggested by an assumption relating only the number of products available to software quality. Failing to account for heterogeneity across firms and software titles precludes any hope of estimating their marginal contribution to a platform.

The second approach – a two-stage procedure – can yield reasonable results only if the econometrician observes all characteristics of consumers onboard each platform that influence their demand for software. Not only are the data requirements far more intensive, but this also rules out the possibility for controlling for selection on any unobservable characteristics. This problem is exacerbated once multiple periods are introduced, since a two-stage estimation procedure also cannot

[39] If there is uncertainty only over the utility derived from the platform itself – and not from software – then demand can be estimated in this manner as well. Such platform-only uncertainty may come from the consumer's own idiosyncratic preferences or the actions of other consumers. E.g., consumers may realize certain needs subsequent to platform adoption or can derive utility from the adoption decisions of other consumers (network effects).

[40] An example of this latter phenomenon would be that consumers do not know exactly the illness(es) that may beset them until after they join an HMO and become sick; consequently, since the utility from a hospital on an HMO plan is contingent on the type of illness acquired, consumers can only form expectations over the utility an HMO's hospital network provides (Ho (2006)). In this setting, the uncertainty will be from the dynamic setting and the durability of the hardware platform: consumers do not know for certain what software titles will be available in the future.

[41] Endogenizing the set of software products onboard each platform is the focus of section 5.

[42] See, e.g., Song and Chintagunta (2003) for PDAs; and Clements and Ohashi (2005), Prieger and Hu (2006), and Corts and Lederman (2007) in videogames.

consistently handle the dynamic evolution and selection of consumer heterogeneity across platforms.

Simultaneous Estimation

The tight integration between hardware and software demand suggests moving towards a method which can simultaneously estimate both sides at once. For expositional purposes, I adopt a discrete choice based approach to demand estimation (see e.g. Lancaster (1971), McFadden (1973), Berry (1994) and Berry, Levinsohn, and Pakes (1995)). These later models can allow consumer preferences for product characteristics to vary as a function of observed and unobserved individual characteristics, and thus allow for reasonable substitution patterns across goods.

Let the total expected lifetime utility that a consumer derives from a single platform be given by

$$U(\psi_i, x_j, \xi_j, \Gamma_j(\cdot); \theta)$$

where ψ_i is a vector of individual characteristics and preferences, x_j and ξ_j are observable and unobservable (to the econometrician) product characteristics respectively (where prices are included in x_j), and $\Gamma_j(\cdot)$, which may also be a function of individual preferences and characteristics, represents the total expected utility a consumer derives from purchasing and using software available on platform j.[43] As mentioned before, so that the choice of optimal bundle of software cannot be collapsed into the hardware purchase decision (which may be conceivable in a static environment, but not so in a dynamic environment studied later), I assume that there is some uncertainty over software quality or availability that is resolved only after the hardware is purchased. θ is a vector of parameters to be estimated, which includes any parameters governing the distribution of unobserved consumer characteristics and preferences.

A consumer will purchase the platform that maximizes her utility, and thus will chose j if and only if:

$$U(\psi_i, x_j, \xi_j, \Gamma_j(\cdot); \theta) \geq U(\psi_i, x_r, \xi_r, \Gamma_r(\cdot); \theta) \ \forall\, r \in J \cup \{0\}$$

where $j = \{0\}$ represents the "outside option" of non purchase. Let

$$A_j = \{\psi : U(\psi_i, x_j, \xi_j, \Gamma_j(\cdot); \theta) \geq U(\psi_i, x_r, \xi_r, \Gamma_r(\cdot); \theta) \ \forall\, r \in J \cup \{0\}\}$$

denote the set of values for ψ which induce consumers to choose good j. If $P_0(d\psi)$ denotes the (initial) population density of ψ, then the share of consumers who choose platform j is given by

$$s_j(x_j, \xi_j, \Gamma(\cdot); \theta) = \int_{\psi \in A_j} P_0(d\psi) \tag{1}$$

With the exception of $\Gamma_j(\cdot)$, the hardware adoption model and aggregation is no different than previous discrete choice demand models. However, as noted earlier, $\Gamma(\cdot)$ is endogenous and comes

[43] I am implicitly assuming total software utility enters linearly into a consumer's calculation of hardware utility, so that the *expected* utility of software $\Gamma_j(\cdot)$ is a sufficient statistic for computing the expected lifetime utility from hardware.

from the software side of the market.

To define $\Gamma(\cdot)$, it becomes necessary to examine the utility a consumer receives from each software title. Consider a consumer who has purchased platform j, and now decides which subsect of software titles $K_j \in \mathbf{K}_j$ to purchase. For expositional purposes, assume that each software product is in an independent market such that there are no substitution or complementarities across titles.[44] Thus, a consumer can decide whether or not to buy a particular title in isolation, and will purchase a given title if that title yields higher utility U^{sw} than the outside good:

$$U^{sw}(\psi_i, w_k, \eta_k; \theta) \geq U^{sw}(\psi_i, w_0, \eta_0; \theta)$$

where now w_k and η_k are the observable and unobservable characteristics of title k. Although the set of consumer types A_k who purchase good k is defined similarly as on the hardware adoption side,

$$A_k = \{\psi : U^{sw}(\psi_i, w_k, \eta_k; \theta) \geq U^{sw}(\psi_i, w_0, \eta_0; \theta)\}$$

the share on platform j of consumers who choose k is given by

$$s_k(w_k, \eta_k; \theta) = \int_{\psi \in A_j \cap A_k} P_0(d\psi) \tag{2}$$

where now the bounds of integration are over the intersection $A_j \cap A_k$. This reflects the fact that consumers who purchased platform j are a selected subsample of the entire population, and it are only these consumers who must be considered when analyzing software demand; this is precisely the selection issue brought up earlier. With independent software titles, software utility on platform j will simply be the maximum utility derived by purchasing or not purchasing each particular software title:

$$\Gamma_j(\cdot) = E[\sum_{k \in \mathbf{K}_j} \max\{U^{sw}(\psi_i, w_k, \eta_k; \theta), U^{sw}(\psi_i, w_0, \eta_0; \theta)\}]$$

with the expectation possibly being over both the set of products $\mathbf{K}_j$ as well as individual product utilities $U^{sw}(\cdot)$.

For a given parameter vector θ, most of the established estimation algorithms require the computation of shares $\{s_j(\cdot;\theta)\}_{\forall j}$ and $\{s_k(\cdot;\theta)\}_{\forall k \in \mathbf{K}_j, \forall j}$: e.g, they are used in the likelihood function for MLE, or are required to recover unobservable characteristics ξ and η from which moments may be constructed (as in Berry, Levinsohn, and Pakes (1995)). But to do so, the aforementioned endogeneity and selection problem must be solved. Without knowing the distribution of consumer types who select a given platform, calculating the share of consumers who purchase a particular software title is impossible; at the same time, computation of the distribution of consumer types who select a given platform requires knowledge of software quality, which itself is predicted from market shares derived from the software adoption side. More formally, for a given θ, the share of

[44] If there were complementarities or substitutability, the model can be extended simply to a consumer choosing the optimal bundle of software titles.

consumers who purchase platform j given by (1) cannot be computed without first knowing $\Gamma_j(\cdot)$; however, computation of Γ_j will require calculation of the utility of each software title on platform j, which in turn requires knowledge of the distribution of consumers who have selected platform j (i.e., the limits of integration in (2)).

This paper introduces a nested fixed point routine to solve this particular issue, discussed in further detail in 3.3. I first fix the distribution of consumer types onboard each platform, and obtain a first-step estimate of the fraction of consumers who purchase each title. This allows for the construction of a first-step estimate of total software quality for each platform. I then update the distribution of consumer types onboard each platform using the new estimated software quality. This procedure is repeated, iterating between estimating hardware adoption (updating the distribution of consumer types) and software adoption (updating the quality of software onboard each platform) until convergence, which occurs when predicted values from the hardware side are consistent with predicted values on the software side.

3.2 Dynamic Model of Consumer Demand

Although the previous discussion illustrated the importance of jointly estimating hardware and software demand and controlling for the selection of heterogenous consumers across platforms, it did so in a static one-period environment. In many applications including the one examined here, this is unrealistic. In the videogame industry, consumers internalize future software availability, quality differences, and potential price drops when deciding when and whether or not to purchase either a hardware system or software title. There is also a clear interdependence of demand across time. First, since videogame consoles and software are durable goods, consumers who have purchased a particular console or title in the past will no consider purchasing the same goods in the future, and the potential market size for a given product will thus shrink over time. Furthermore, the types of consumers who purchase earlier will be different that those who purchase later: just as there is selection across platforms, there is also selection across time – e.g., early adopters will tend to be less price sensitive or demonstrate a higher degree of affinity for these products. Controlling for these dynamic issues is crucial in accurately estimating demand in this and most platform environments.

I begin by modifying the notation sofar introduced to be time-specific. Now, in each period t, there are J_t hardware consoles and $K_{j,t}$ software titles on each platform available for purchase. Each consumer who has not yet purchased a console before time t may choose to do so, or wait until the next period. In addition, any consumer who has purchased a console by time t may decide to purchase any software title $k \in \mathbf{K}_{j,t}$ she has not yet already purchased, where $\mathbf{K}_{j,t}$ represents the set of software titles available on platform j at time t. Since hardware and software are durable goods, and consumers have expectations over the evolution of these goods' qualities and prices, the timing of purchase becomes a dynamic optimization problem – not only which product should a consumer purchase, but also *when* should she purchase.

I will first describe the consumer hardware adoption decision before moving on to discussing software adoption. For expositional purposes, I introduce a model whereby consumers may only purchase one hardware console before extending it to allow consumers to purchase multiple consoles.

Hardware Adoption

At any period of time, consumers who have not yet purchased a hardware platform may choose to do so. The total lifetime (expected) utility of consumer i who purchases platform j at time t is given by:

$$u_{ijt} = \underbrace{\alpha_i^x x_{j,t} - \alpha_i^p \ln(p_{j,t}) + \Gamma_{j,t}(\alpha_i^p, \alpha_i^\gamma) + \xi_{j,t}}_{\delta_{ijt}} + \epsilon_{ijt} \tag{3}$$

where $x_{j,t}$ are observable characteristics of platform j at time t (which include a platform-specific and monthly fixed effects, age, age squared, and the current platform installed base), $p_{j,t}$ is the price of the console, $\Gamma_{j,t}$ is the expected present-discounted value of being able to purchase software for the platform in the current and future periods, $\xi_{j,t}$ is a product characteristic observable to the consumer but not to the econometrician, and ϵ_{ijt} is an individual-platform-time specific component which represents idiosyncratic consumer heterogeneity unobservable to the econometrician but realized by the consumer only at time t. Additionally, $\{\alpha^x, \alpha_i^p, \alpha_i^\gamma\}$ are (possibly individual specific) coefficients which reflect how intensely a consumer prefers platform characteristics, price, and software. The actual functional form of $\Gamma_{j,t}(\cdot)$ emerges from the software adoption portion of the model which will be described in the next subsection; the only restriction made on Γ here is that it differs across agents only as a function of their income and software preference coefficient $\{\alpha_i^p, \alpha_i^\gamma\}$, and enters linearly into the utility specification. Finally, I denote the portion of product's individual specific utility net the individual-specific unobservable by $\delta_{i,j,t}$; it can be thought of as the price-adjusted quality for platform j, and if $\epsilon_{i,j,t}$ were mean zero, it would represent the *mean utility* of such a purchase. Instead of purchasing a console, a consumer also may choose to wait and consume the outside good for one period – which yields utility $u_{i0t} = \epsilon_{i0t}$ – and return to the market in the next period.

In each period, a consumer chooses her optimal action – buy today or wait until next period – given her preferences, current product qualities, prices and software availability, and expectations over the evolution of these characteristics. A consumer's utility from being on the market for a hardware platform – contingent on following the optimal policy – is given by the following value function:

$$V_i(\epsilon_{i,t}, \Omega_{i,t}) = \max\{\max_j u_{ijt}, u_{i,0,t} + \beta E[V_i(\epsilon_{i,t+1}, \Omega_{i,t+1}]\} \tag{4}$$

where $\Omega_{i,t}$ includes current product attributes as well as the time of year (in order to account for seasonality effects) and any other market characteristics which may affect firm product pricing, entry, exit, or attributes such as installed base. In general, it includes all variables at time t in the consumer's information set which affect her utility or value for waiting. This value function defines an optimal stopping problem which specifies when a consumer should (if ever) purchase a console.

For tractability, I first impose the following assumption on the idiosyncratic utility shocks:

Assumption 3.1. $\{\epsilon_{i,t}\}_{\forall i,t}$ *are independently and identically distributed type 1 extreme value with variance normalized to* $\pi^2/6$.

Whereas the normalization of the variance pins down the scale of utility (which is itself not identified), the distributional and conditional independence assumption allows (4) to be analytically integrated over ϵ and provide an "expected" value function (EV) for consumer i as a function of current state variables:

$$EV_i(\Omega_{i,t}) = \int_{\epsilon_{i,t}} V_i(\epsilon_{i,t}, \Omega_{i,t}) dP_\epsilon = \ln(\exp\{\delta_{it}\} + \exp\{\beta E[EV_i(\Omega_{i,t+1})]\}) \tag{5}$$

where

$$\delta_{i,t} \equiv E_{\varepsilon_t}\{\max_j u_{i,j,t}\} = \ln(\sum_{j \in J_t} \exp(\delta_{i,j,t}))$$

and represents the expected utility from purchasing a product which delivers the maximal utility at time t.[45] Also known as the "inclusive value" for consumer i, $\delta_{i,t}$ is a sufficient statistic for determining if the consumer will participate in the market at time t by purchasing any hardware platform.

Despite reframing the value function in expectations as in (5), the state space is still too large for the consumer's optimization problem to be computationally solvable. One solution, utilized by Melnikov (2001), Hendel and Nevo (2006) and Gowrisankaran and Rysman (2007), is to assume that the inclusive values ($\delta_{i,t}$) within a market follow a first-order Markov process. However, although significantly reducing the computational burden of estimation, this assumption nonetheless imposes strong restrictions on the nature of industry competition and evolution.[46] Rather than follow this approach, I instead assume that consumers perceive the mean utilities $\delta_{i,j,t}$ for each console to evolve according to an exogenous first-order process, and depends on previous values of itself in addition to $\{\delta_{i,j',t}\}_{j' \neq j}$ of *all other competing hardware platforms*, as well as the time of year:

Assumption 3.2. *(Consumers perceive that)* $\{\delta_{i,j,t}\}_{j \in J_t}$ *can be summarized by an exogenous first-order Markov Process:*

$$F(\{\delta_{i,j,t+1}\}_{j \in J_{t+1}} | \Omega_{i,t}) = F_i(\{\delta_{i,j,t+1}\}_{j \in J_{t+1}} | \{\delta_{i,j,t}\}_{j \in J_t}, m(t)) \tag{6}$$

where $m(t)$ *represents the month at time* t.

In addition to allowing each individual i to have different expectations over the evolution of the industry (F_i being individual specific), this specification is more general by allowing each product's "quality" to evolve and as a function of – among other things – the proximity of other products' quality to its own. Furthermore, which is crucial for this industry, any monthly seasonality is

[45] Ben-Akiva (1973) shows that $E_\epsilon[\max_i\{x_i + \epsilon_i\}|x_i] = \ln(\sum_i \exp(x_i))$ when $\{\epsilon_i\}$ are i.i.d. extreme value type 1. Application of this twice provides (5). See Rust (1987) and Melnikov (2001) for use of this in a similar context.

[46] E.g., the inclusive value for a consumer $\delta_{i,t}$ may be high if there are several products with low prices or relatively few products with high prices; a consumer would not only exhibit the same probability of purchasing but also have the same expectation over future values of $\delta_{i,t}$ in both cases.

captured as a state variable. However, this assumption is still problematic in that it is difficult to create a supply model which generates first-order processes in $\delta_{i,j,t}$, which includes prices and software title availability.[47] At the same time, as noted in Hendel and Nevo (2006), such first order processes are reasonable approximations to consumer expectations and memory as consumers may remember prices and software availability from only their previous visit. Higher order processes not only may be too burdensome for estimation, but for consumer decision making as well.

Combining equations (5) with assumption (6), the consumer's expected value function can now be rewritten as:

$$EV_i(\{\delta_{i,j,t}\}_{j\in J_t}, m(t)) = \ln(\exp(\delta_{i,t}) + \exp(\beta E[EV_i(\{\delta_{i,j,t+1}\}_{j\in J_t}, m(t+1))|\{\delta_{i,j,t}\}_{j\in J_t}, m(t)])) \quad (7)$$

where now the state space has been drastically reduced from $|\Omega_{i,t}|$ to one of only at most $J+1$-dimensions. Due to the limited number of platforms in the videogame industry (only 3 in the time period analyzed), this state space is now small enough for implementation.

Software Adoption

I now turn to analyze the software purchase decisions for a consumer, used to construct the "software quality" function $\{\Gamma_{j,t}(\cdot)\}_{\forall j,t}$ in (3).

In each period t, a consumer who has purchased platform j in any period $\tau \leq t$ enters the market and may purchase any software title $k \in \mathbf{K}_{j,t}$ she has not yet purchased. Crucially, I assume each consumer views each software title as a separate market – i.e., the decision to purchase a title k is independent of purchasing $k' \neq k$. Although this assumption is motivated by feasibility,[48] it does not seem to be overly restrictive in this particular industry for reasons discussed earlier. For expositional purposes, the remainder of this subsection omits the j subscript for the platform and, unless otherwise specified, values are assumed to be platform specific.

A consumer's expected lifetime utility from buying title k in period t (provided she already owns the platform) is given by:

$$v_{i,k,t} = \tilde{\alpha}_i^{\gamma} + \tilde{\alpha}^w w_{k,t} + \tilde{\eta}_{k,t} - \alpha_i^p \ln(p_{k,t}) + \tilde{\epsilon}_{ikt} \quad (8)$$

where $w_{k,t}$ are observable software characteristics (which include a game-specific fixed effect, monthly fixed effects, as well as age, age squared, and the current installed base of previous purchasers), $\tilde{\eta}_{k,t}$ is an software characteristic unobservable to the econometrician (but observable to the consumer), $p_{k,t}$ the price, and $\tilde{\epsilon}_{ikt}$ is an individual-software-time specific utility shock. $\tilde{\alpha}_i^{\gamma}$ is an individual spe-

[47] Setting aside dynamic issues due to the evolving nature of the installed and remaining customer base, there is a strong parallel to the industry dynamics model introduced in Pakes and McGuire (1994) and Ericson and Pakes (1995): if $\delta_{i,j,t}$ represented platform j's "index of efficiency" and platforms could engage in investment efforts to improve this value, then under certain assumptions (including the absence of other strategic control variables) there would exist a Markov-Perfect Nash equilibrium whereby $\{\delta_{i,j,t}\}_{\forall j}$ would evolve according to a first-order transition kernel.

[48] In a dynamic environment, individually tracking each consumer's inventory and subsequent choice set is too computationally burdensome. Such problems were absent in the static environment used for exposition in Section II.

cific preference for "gaming" reflected in the increase in utility of any particular piece of software, and α_i^p reflects price sensitivity, and importantly is the same coefficient as on the hardware side. A consumer can also decide not to buy a piece of software at time t and return to the market in the next period, yielding the outside option utility $v_{ik_0t} = \tilde{\epsilon}_{ik_0t}$.

Mirroring the hardware side, I make a distributional assumption on the individual-specific utility shocks:

Assumption 3.3. *$\{\tilde{\epsilon}_{i,k,t}, \tilde{\epsilon}_{i,k_0,t}\}_{\forall i,k,t}$ are independently and identically distributed extreme value type 1 with variance $\sigma_\epsilon^2\pi^2/6$.*

Here I allow for the variance of unobserved heterogeneity – σ_ϵ – to vary between the hardware and software sides. I must account for this when combining measures of utility across sides as any shared coefficients (e.g., α_i^p) need to be appropriately scaled.[49] I thus re-express the software utility in (8) by multiplying and dividing through by σ_ϵ:

$$v_{i,k,t} = \sigma_\epsilon(\underbrace{\alpha_i^\gamma + \alpha^w w_{k,t} + \eta_{k,t} - \alpha_i^{p,sw}\ln(p_{k,t})}_{\zeta_{ikt}} + \epsilon_{ikt})$$

where $\{\alpha_i^\gamma, \alpha^w, \alpha_i^{p,sw}, \eta_{k,t}, \epsilon_{i,k,t}\} = \{\tilde{\alpha}_i^\gamma, \tilde{\alpha}^w, \alpha_i^p, \tilde{\eta}_{k,t}, \tilde{\epsilon}_{i,k,t}\}/\sigma_\epsilon$, and $\zeta_{i,k,t}$ represents the (scaled) utility of purchasing a piece of software net of individual-specific-unobservable $\epsilon_{i,k,t}$, and may also be referred to as the price-adjusted quality or mean-utility for software k. To prevent confusion, I will use $\alpha_i^{p,hw} \equiv \alpha_i^p$ to refer to the coefficient on price used on the hardware side whenever appropriate.

A consumer's optimal stopping problem for when (if ever) to purchase software title k is given by:

$$W_i(\Omega_{k,t}, \epsilon_{ikt}) = \max\{v_{ikt}(\omega_{k,t}, v_{ik_0t} + \beta E[W_i(\Omega_{k,t+1}, \epsilon_{i,k,t+1})]\}$$

where now $\Omega_{k,t}$ represents any relevant variables which influence consumer i's utility from purchasing or waiting for title k. Again, to reduce the dimensionality of the state space, the following assumption is made on the evolution of each software-title's mean-utility:

Assumption 3.4. *(Consumers perceive that) ζ_{ikt} can be summarized by an exogenous first-order Markov process:*

$$G(\zeta_{i,k,t+1}|\Omega_{k,t}) = G_{i,j}(\zeta_{i,k,t+1}|\zeta_{ikt}, m(t)) \tag{9}$$

where now $G_{i,j}$ is allowed to be specific to individual i and console j. It is, however, not title specific, and thus all software titles on a given platform are perceived by consumers to follow the same evolutionary path contingent on their price-adjusted quality and time of year. This assumption is subject to the same caveats and support as provided earlier.

Given assumptions 3.3 and 3.4, the expected value function of being in market for software title k at time t can be re-expressed as

$$EW_i(\zeta_{ikt}) = \int_{\epsilon_{i,k,t}} W_i(\zeta_{i,k,t}, \epsilon_{i,k,t})dP_\epsilon = \sigma_\epsilon \ln(\exp(\zeta_{ikt}) + \exp(\beta E[EW_i(\zeta_{ikt+1})|\zeta_{ikt}]))$$

[49] C.f. Train (2003), chapter 2 for further discussion.

which imbeds consumer i's expectations over future prices and characteristics for software k.

To close the software adoption side, I need to link $\Gamma_{j,t}$ to the value of being on the market for software on platform j. This "total software utility" on platform j can be separated into two parts: (i) the utility from software available in the present period, and (ii) the utility from new software that will arrive in future periods. Let this latter value be denoted $\Lambda_{j,t}(\alpha_i^p, \alpha_i^\gamma)$, and let $\mathbf{K}_{j,t}^R$ denote the set of software titles released on platform j at time t. Then $\Gamma_{j,t}$ is given by:

$$\Gamma_{j,t}(\alpha_i^p, \alpha_i^\gamma) = \underbrace{\left[\sum_{k \in \{\bigcup_{\tau \le t} \mathbf{K}_{j,\tau}^R\}} EW_{ikt}(\zeta_{i,k,t}) \right]}_{\text{Current Software Utility}} + \underbrace{\Lambda_{j,t}(\alpha_i^p, \alpha_i^\gamma)}_{\text{(Expected) Future Software Utility}} \tag{10}$$

where the first term aggregates the expected utility of being on the market for each piece of software currently available; it imbeds the consumer's optimal policy in its calculation.

Although all three consoles survived to the end of observed time period, other consoles in previous generations have prematurely "died" and left the market. To allow for this uncertainty, I allow consumers to believe a console will continue to have software released in the next period with probability β_γ, constant across time and consoles. Conditional on β_γ, if a consumer had perfect information over all future titles that would be released in addition to the terminal date T, future utility would be specified by

$$\tilde{\Lambda}_{i,j,t} = \sum_{\tau=1}^{T} (\beta_\gamma \times \beta)^\tau \left(\sum_{k \in \mathbf{K}_{j,t+\tau}^R} EW_{i,k,t+\tau}(\zeta_{i,k,t+\tau}) \right) \tag{11}$$

A consumer does not know exactly the number nor quality of future titles, however. I deal with this by assuming consumers have rational expectations consistent with the observed data, and condition on current observed market variables and product characteristics; i.e., $\Lambda_{i,j,t} = E[\tilde{\Lambda}_{i,j,t} | \Omega_t, \beta_\gamma]$. In estimation, as will be described in the next section, I will use a nonparametric series regression of (11) on console characteristics and software availability as an approximation to proxy for a consumer's expectations.

In the appendix, I discuss an alternative formulation for future software utility whereby consumers believe the number and quality of titles released in each period are drawn independently from estimated distributions consistent with the data.

Multiple Hardware Purchases

When consumers are allowed to multihome and purchase multiple consoles, both the dynamic optimization problem as well as the expected lifetime utility from hardware purchase changes. First, a consumer upon purchasing a console no longer leaves the market, but retains the option value of returning in a future period and acquiring a second or even third console. Second, the value of a console is different if that console is the first or second (or third) purchased: the expected utility derived from purchasing an additional console should not include utility from a software title the user already had access to on her original console; there may also be an additional impact on hardware utility due to the existence of complementarities or substitution between newly purchased and already owned consoles).

To account for multiple hardware purchases, let there now be a new state variable indicating consumer i's inventory of hardware consoles she already owns at time t, and represent this by $I_{i,t} \in \mathbf{I} \equiv \{0,1\}^3$. The expected lifetime utility from purchasing a new console now becomes a function of which consoles a consumer already owns, and is represented by:

$$u_{ijt}(I_{i,t}) = \underbrace{\alpha_i^x x_{j,t} - \alpha_i^p \ln(p_{j,t}) + \Gamma_{j,t}(\alpha_i^p, \alpha_i^\gamma; I_{i,t}) + D(I_{i,t}) + \xi_{j,t}}_{\delta_{ijt}(I_{i,t})} + \epsilon_{ijt}$$

where software utility $\Gamma_{j,t}$ is now a function of $I_{i,t}$ and is adjusted to account for the fact that a user may have already had access to certain titles on the consoles she already owns: $\Gamma_{j,t}$ is defined as in (10), except now only includes those titles $k \in \mathbf{K}_j \setminus \mathbf{K}_{j' \forall j' \in I_{i,t}}$.

I also introduce a new term $D(I_{i,t})$ which reflects the complementarity or substitutability effect(s) that may exist with ownership of multiple consoles.[50] Conceptually this term should be negative, as the utility of a console would seem to decrease if another is already owned. However, since much of this diminished utility comes from the fact that non-exclusive software titles on additional platforms were already accessible if another console was already owned, any remaining effects should come only from other factors such as time constraints or saturation of gaming needs.[51]

In the appendix, I discuss how to modify a consumer's value function (7) and assumption 3.2 on beliefs in order to integrate this new utility function into the dynamic model.

3.3 Estimation and Computation

Berry, Levinsohn, and Pakes (1995) and following literature typically estimates discrete choice models by recovering the set of unobserved product characteristics for any parameter vector θ which perfectly rationalize the model's predicted market shares with observed market shares, and then using a generalized methods of moments (GMM) estimator based on forming conditional moments with these unobserved characteristics: i.e., the identifying condition is typically $E[\xi|z] = 0$, where z is a set of instruments orthogonal to ξ.[52] However, this process requires finding appropriate instruments for observed characteristics including price, which may not exist.

Rather than proceed in this fashion, I instead leverage the dynamic aspect of my data and estimate based on the predicted evolution of the unobserved product characteristics. Namely, I assume that ξ and η for each hardware system and software title evolve according to an exogenous Markov process where the changes are independent from *changes* in observed characteristics.[53] In

[50] This is similar to the additive complementarity or substitution term used in Gentzkow (2007).

[51] These factors are not estimated to be significant under the fully specified model when D is a constant.

[52] Other applications include Nevo (2001) and Petrin (2002).

[53] Berry, Levinsohn, and Pakes (1995) (in section 4) notes the possibility of proceeding in this fashion when utilizing panel data. In simultaneous work, Sweeting (2007) uses a similar assumption on the evolution of unobserved

turn, I estimate via (conditional) maximum likelihood (ML) on the probability of observing the changes in the error terms.

Formally, the econometrician observes the characteristics and quantities of each console and software title in every period. Let r_j denote the release date for console j and r_k the release date for software title k. From the model, the implied values of the unobserved product characteristics $\{\xi_{j,t}\}_{1\leq j\leq J, r_j\leq t\leq T}$ and $\{\eta_{j,k,t}\}_{1\leq j\leq J, r_k\leq t\leq T, k\in K_{j,t}}$ can be computed to rationalize observed with predicted market shares as a function of parameters θ to be estimated. I assume the following:

Assumption 3.5. *Unobserved product characteristics for each console and software title evolve according to a first-order autoregressive (AR(1)) process, where the errors*

$$\begin{aligned}
\nu^{hw}_{j,t}(\theta) &= \xi_{j,t}(\theta) - \rho^{hw}\xi_{j,t-1}(\theta) \qquad \forall j \in J_t \\
\nu^{sw}_{j,k,t}(\theta) &= \eta_{j,k,t}(\theta) - \rho^{sw}\eta_{j,k,t-1}(\theta) \qquad \forall j \in J_t,\ \forall\, k \in K_{j,t-1}
\end{aligned} \tag{12}$$

are independent of each other and changes in all observed characteristics, and are identically distributed according to probability densities $f^{hw}(\cdot;\theta)$ and $f^{sw}(\cdot;\theta)$, respectively.[54]

The likelihood of observing a set of these errors is given by:

$$L(\theta) = \prod_{j=1}^{J}\left(\prod_{t=r_j+1}^{T} f^{hw}(\nu^{hw}_{j,t}(\theta);\theta) \times \prod_{t=r_k+1}^{T}\prod_{k=1}^{K_{j,t}} f^{sw}(\nu^{sw}_{j,k,t}(\theta);\theta)\right)$$

and the log-likelihood by:

$$\mathcal{L}(\theta) = \sum_{j=1}^{J}\left(\sum_{t=r_j+1}^{T} \ln f^{hw}(\nu^{hw}_{j,t}(\theta);\theta) + \sum_{t=r_k+1}^{T}\sum_{k=1}^{K_{j,t}} \ln f^{sw}(\nu^{sw}_{j,k,t}(\theta);\theta)\right) \tag{13}$$

The estimate of θ_0 is the θ that maximizes (13):

$$\hat{\theta} = sup_{\theta\in\Theta}\mathcal{L}(\theta) \tag{14}$$

The initial values for $\{\xi_{j,r_j}\}_{\forall j}$and $\{\eta_{j,k,r_k}\}_{\forall j,k}$ are conditioned on in the specification of the likelihood. However, as T grows large, the contribution of these initial values to the likelihood grow negligible. Provided $|\rho| < 1$, the exact ML estimator and this conditional ML estimator which omits the first ξ_{j,r_j} and η_{j,k,r_k} term will have the same asymptotic distribution.[55] Leveraging the dynamic aspects of the problem in this fashion not only allows us to condition for an initial conditions problem,[56] but it proves robust to the possibility that hardware and software release

characteristics within a GMM estimator.

[54]The stationarity coefficients $(\rho^{hw}, \rho^{sw}) \in \theta$ and need to be estimated as well. Since the drift of a MA(1) process is not separately identified from the level of product fixed effects contained within the α^x and α^y coefficients, it is fixed to be 0.

[55]See e.g. Fuller (1996); the estimator is also similar to the "y_0-conditional estimator" described in section 8.7 of Hayashi (2000).

[56]I.e., the initial values of ξ or η upon release may be correlated with product observable characteristics.

dates are "timed": e.g., titles that have a relatively high initial unobserved quality η_{j,k,r_k} may systematically be released during "high" demand months, such as the beginning of summer and during the holiday season.[57] As long as its evolution continues to follow (12), any such strategic timing will not bias estimates.

The following proposition proves this estimator is consistent and asymptotically normal.

Proposition 3.6. *Let $n_j^{hw} = \sum_{t=r_j}^{T} 1$ and $n_j^{sw} = \sum_{t=r_k}^{T} \sum_{k=1}^{K_{j,t}} 1$ be the number of observations for platform j, and $N = \sum_{j=1}^{J}(n_j^{hw} + n_j^{sw})$ be the total number of error observations. Provided certain identifying and invertibility assumptions are satisfied (see assumption C.1 in Appendix), if $\hat{\theta}$ is a solution to* (14)*, then as $T \to \infty$, $\mu_j^{hw} = n_j^{hw}/N$ and $\mu_j^{sw} = n_j^{sw}/N$ constant,*

I. $\hat{\theta} \to_p \theta_0$

II. $\sqrt{N}(\hat{\theta} - \theta_0) \to_d N(0, \mathbf{J}_0^{-1})$*, where*

$$\mathbf{J}_0 = \left[\sum_{j=1}^{J} \mu_j^{hw} E\left(\frac{\partial \ln f^{hw}}{\partial \theta_r} \frac{\partial \ln f^{hw}}{\partial \theta_s} \right)\Bigg|_{\theta=\theta_0} + \mu_j^{sw} E\left(\frac{\partial \ln f^{sw}}{\partial \theta_r} \frac{\partial \ln f^{sw}}{\partial \theta_s} \right)\Bigg|_{\theta=\theta_0} \right] \tag{15}$$

Proof. See Appendix. □

The proof requires modifying the standard asymptotic arguments for maximum likelihood estimators (e.g., Rao (1973)) by accounting for the non-identical distribution of the hardware and software errors in addition to the varying number of observations across platforms and time. For estimation, I will assume that the errors are normally distributed with $\nu_{j,t}^{hw} \sim N(0, \sigma_{hw}^2)$ and $\nu_{j,k,t}^{sw} \sim N(0, \sigma_{sw}^2)$.

Note also ξ and η cannot be computed exactly, but rather only approximately, due to both population sampling error and the need to simulate the integrals defining market shares for each product.[58] As in Berry, Levinsohn, and Pakes (1995), I assume the population sampling is negligible due to the large sample size of over 100M U.S. households. However, the issues introduced with the need for simulation cannot be ignored. As has well been documented, simulation error is problematic in these contexts when market shares are small (Berry, Linton, and Pakes (2004)), and it also introduces a bias due to the non-linearity of the log transformation in the maximum likelihood estimator (Lee (1995)). To avoid the problems involved with simulation error, I instead discretize the distributions of consumer heterogeneity and assume consumers belong to one of several "types" of consumers with identical $\{\alpha_i^\gamma, \alpha_i^p\}$ coefficients (but with different realizations of $\{\epsilon_{it}\}$).

To calculate standard errors, I use the estimate $\hat{\theta}$ and take the sample analogue of (15).

Parameters to Estimate

Let $\theta_1 = \{\rho^{hw}, \rho^{sw}, \tilde{\alpha}_0^{p,sw}, \sigma_{p,sw}, \sigma_\gamma, \sigma_\epsilon, \beta_\gamma, D\}$ and $\theta_2 = \{\alpha^x, \alpha^w\}$. The parameters to be estimated are $\theta = \{\theta_1, \theta_2\}$.

[57] See Einav (2007) for issues regarding seasonality and timing in the US Motion Picture Industry.

[58] These are provided by equations (26) and (31) in the appendix.

Computation

Here I overview the procedure of recovering the unobserved product characteristics $\{\xi_{j,t}(\cdot)\}_{\forall j,t}$ and $\{\eta_{j,k,t}(\cdot)\}_{\forall j,k,t}$ as a function of the parameter vector θ. Once these values are obtained, the log-likelihood function in (13) can be computed.

The approach for recovering the unobservable utility components ξ in the hardware side and η in the software side is similar to the approaches utilized in Nair (2007) and Gowrisankaran and Rysman (2007), which both in turn nest the methodologies of Rust (1987), Berry (1994), and Berry, Levinsohn, and Pakes (1995). This paper, however, links these two demand systems and estimates both sides simultaneously; additionally, in allowing each platform's mean utility to evolve separately from the inclusive value of the industry and explicitly accounting for seasonality effects, the specification here is more general.

For any given θ parameter, we first assume starting values for $\{\Gamma_{j,t}\}_{\forall j,t}$, which can either be fixed (e.g., at 0), or can be computed by assuming that the distribution of consumer heterogeneity across new console purchasers is stationary and first estimating the software side. I employ the latter approach. Utilizing these initial values for Γ^0, I first proceed with estimating the hardware adoption side. The "mean" utility $\delta_{j,t}$ for each platform that rationalizes predicted market shares to observed market shares is found via an iterative contraction mapping introduced in Berry, Levinsohn, and Pakes (1995). For each iteration of this contraction mapping, each consumer's belief process over the evolution of $\delta_{i,j,t}$ for every inventory state ι is updated according to the regression

$$\begin{aligned} &F_i(\delta_{i,j,t+1}(\iota)|\{\delta_{i,j,t}(\iota)\}_{\forall j}, \iota, m(t)) = \\ &\qquad \varphi_{i,j,\iota,0} + \sum_{j'=1}^{3} \varphi_{i,j',\iota,1}\delta_{i,j',t}(\iota) + \sum_{m=1}^{11} \varphi_{i,j,\iota,m+3}\chi_m(t) + \upsilon_{i,j,\iota,t} \end{aligned} \tag{16}$$

(whereby $\chi_m(t)$ are indicator variables if t is in month m) and her optimal stopping problem is solved to determine the probability of purchase.[59] At each point in time, the number and identity of consumers at each inventory state evolves according to (30), which at the end is aggregated across consumers to form predicted market shares for each month.

Once the hardware adoption side is solved for a given vector of $\{\Gamma_{j,t}(\iota)\}_{\forall j,t,\iota}$, I use the probability that each consumer adopts a hardware platform in each period to form the consumer distribution of each hardware's installed base, denoted by $\{dP_{j,t}(\alpha^p, \alpha^\gamma)\}_{\forall j,t}$. This updated distribution of consumer types is then used to estimate the software adoption decision, which proceeds via a similar nested framework as the hardware side except that it now must be done for each console separately. I.e., the same "δ"-contraction mapping is used to recover the mean utilities $\zeta_{j,k,t}$ for each piece of software on a given console, whereby in each iteration consumer expectations are

[59] Unlike using more lagged terms, which would increase the state space and be too computationally expensive, the functional form can be expanded to utilize higher order terms and/or interactions between $\delta_{j,t}$ and its competitors $\{\delta_{j',t}\}_{\forall j' \neq j}$.

updated according to

$$G_{ij}(\zeta_{i,k,t+1}|\zeta_{i,k,t}) = \varphi_{i0}^{j} + \varphi_{i,1}^{j}\zeta_{i,k,t} + \varphi_{i,2}^{j}(\zeta_{i,k,t})^2 + \sum_{m=1}^{11}\varphi_{i,m+2}^{j}\chi_m(t) + \upsilon_{i,k,t}^{j} \tag{17}$$

and the consumer's optimal stopping problem is solved. After the ζ's converge for a given probability distribution $dP(\alpha^p, \alpha^\gamma)$, an updated value of $\Gamma_{j,t}(\cdot)$, given by (10), is computed for every inventory state. Future utility in $\Gamma_{j,t}(\cdot)$ is obtained via a non-parametric series regression on (11) using third-order terms and a full set of interactions on a console's age, number of active software titles, and month dummies.

Finally, this updated value of $\Gamma_{j,t}(\cdot)$ is fed back and the hardware adoption side is re-estimated. The procedure iterates between estimating the hardware and software sides until $\Gamma_{j,t}(\cdot)$ and $P_{j,t}$ converge, at which point ξ and η can be recovered from the final computed values $\delta_{j,t}$ and $\zeta_{j,k,t}$ via a linear regression (at which point θ_2 can be "concentrated" out). A non-derivative based Nelder and Mead (1965) simplex algorithm is used to search for $\hat{\theta}_1$.[60]

The procedure is illustrated in figure 2 with further details provided in Appendix A.3. With an appropriately fine grid and large enough bounds on the state spaces and following an initial calibration period, no problems with convergence for $\delta_{j,t}$, $\zeta_{j,k,t}$, EV_i, EW_i, or $\Gamma_{j,t}$ were encountered during estimation.

3.4 Identification

As in estimation of most dynamic processes, without further restrictions on the discount factor and parameterization of consumer heterogeneity, the model remains unidentified (Rust (1994), Magnac and Thesmar (2002)). I fix the discount rate β at .99. To parameterize consumer heterogeneity, price sensitivity for hardware and software takes the form $\alpha_i^{p,\cdot} = \alpha_0^{p,\cdot} - \sigma_{p,\cdot} y_i$ where y_i is consumer i's annual household income, and $\alpha_0^{p,\cdot}$ and $\sigma_{p,\cdot}$ are parameters to be estimated. As in Berry, Levinsohn, and Pakes (1995), disposable household income y_i for the population is assumed to be (independently) distributed log normally with mean and standard deviation estimated separately from the March 2001 Current Population Survey (CPS) and draw from this distribution.[61] Consumer preferences for software α^γ is assumed independently distributed normally with standard deviation σ_γ; since α^γ enters linearly in utility, its mean is not separately identified from shifts in each software title's fixed effect and hence is normalized to 0.

Even with these restrictions, there still remains the question of what precisely identifies the elements of θ. Unlike in static environments, the dynamic panel context considered here allows for

[60] As shown in the appendix, the non-linear search is only over θ_1 as θ_2 can be "concentrated" out as a function of θ_1 – i.e., the likelihood to be evaluated can be expressed as $\mathcal{L}(\theta_1, \theta_2(\theta_1))$.

[61] This data reports the distribution of disposable household income in 2000, which corresponds to the start of the data. From the Census Bureau, disposable income is defined as including money income, the value of noncash transfers (food stamps, public or subsidized housing, and free or reduced-price school lunches), imputed realized capital gains and losses, and imputed rate of return on home equity. It deducts imputed work expenses, federal payroll taxes, federal and state income taxes, and property taxes on owner-occupied homes.

the repeated observations across time of product characteristics for both hardware and software. Consequently, the components of α^x and α^w – which include product and month level fixed effects as well as age, age squared, and installed base terms – are identified from the time variation in sales as these characteristics change.[62] In a similar fashion, σ_ϵ (the ratio between software utility and hardware utility, which affects how $\Gamma_{j,t}$ enters into hardware demand) and β_γ (the additional discounting assigned to future games that have not yet been released) are identified from changes in hardware demand in response to the release of current and future titles.

The identification of the mean and variance of the consumer heterogeneity parameters $\alpha_0^{p,\cdot}, \sigma_{p,\cdot}, \sigma_\gamma$ in addition to the complementarity effect $D(I_{i,t})$ requires greater discussion, especially in the absence of micro-level data. First, mean household price sensitivities $\alpha_0^{p,\cdot}$ are identified via monthly variation in prices. Since there is very little variation in hardware prices over time (only 2 significant price drops for each console) but far more in software prices, I estimate $\alpha_0^{p,sw}$ and leverage the restriction imposed by the model that $\alpha_0^{p,hw} = \sigma_\epsilon \alpha_0^{p,sw}$. I find that estimating α_0^p and $\alpha_0^{p,sw}$ separately cannot reject this restriction, and thus only report estimates from the latter specification.[63]

Typically, the variance in consumer preferences $\{\sigma_{p,\cdot}, \sigma_\gamma\}$ can be identified from variation in product characteristics: as characteristics change for one product, substitution to products with similar characteristics indicate the presence of heterogeneity; on the other hand, if consumers substitute equally to and from all goods, then consumers are likely to be more homogenous in their preferences. Unfortunately, this argument has limited power when there are only 3 hardware platforms in the data, and substitution away from a particular software title is limited to only the outside good (i.e., waiting until the next period).

Nonetheless, the dynamic nature of the panel data provides another means of identification for consumer heterogeneity: the endogenous shift in the distribution of consumer valuations over time. In this model, consumers who are less price sensitive or have a higher intensity for gaming will select out and adopt a platform earlier than others, creating a difference in the composition of non-adopters and adopters of a hardware system. If household heterogeneity in either α^p and α^γ is substantial, then consumer responses over time to changes in price or software availability *on* a given platform versus *for* a given platform will be different. For example, in the absence of heterogeneity in α^γ, the release of a software title early in the life of a console that attracts 50% of a console's installed base should have the same effect on inducing new consumers to purchase that console as a software title released a year later that attracts the same share: if there was no selection across time and consumers were relatively homogenous, early adopters have the same preferences as later adopters and the predicted software title fixed effect for both games will be the same. However, in the presence of heterogeneity, the installed base of a console will have a higher share of consumers with a high value of α^γ earlier than later. As such, a title released later in a console's history that attracts the same share of consumers as a title released earlier is actually a "higher" quality title

[62] Although not necessary, assuming age effects are shared across hardware platforms and across software titles further helps aid in identification. I assume monthly seasonality effects are the same across years and shared across hardware platforms, but may differ across platforms for software demand.

[63] I also scale the variance of hardware price sensitivity in a similar manner: i.e., $\sigma_{p,hw} = \sigma_\epsilon \sigma_{p,sw}$.

(as reflected by a higher title-fixed effect), since being released later means it must have appealed to a less predisposed base of users; consequently, two titles released at different times with similar demand onboard a platform but with differential effects on total platform demand indicates the presence of heterogeneity in gaming preferences. Similarly, heterogeneity in price sensitivity implies that consumers onboard a console should be growing more price sensitive to software over time as more and more price sensitive consumers purchase hardware.

Finally, I discuss the identification of the complementarity term $D(I_{i,t})$. In estimation, I assume $D(\cdot) = D$ if a consumer already owns at least one other console, and $D(\cdot) = 0$ otherwise. Although possible to be more general, estimation of a single constant is sufficient to capture the possibility that a previously purchased console reduces the future utility from another device *beyond that* of reducing the number of unique titles that can be accessed.[64] Since I fix the total potential market size for video game consoles and assume it to be observable and equal to the number of television households, identification of D comes immediately from its impact on the rate of change in potential consumers for hardware. If D is extremely high, then any purchaser of a hardware system essentially "leaves" the market and does not purchase another console; when D is low, previous purchase does not remove that consumer from consideration when purchasing other hardware devices. As the potential market size directly enters into calculation of the observed and expected share of consumers who purchase products, the rate at which these shares change (and its implied impact on predicted product unobservables) identifies D. It is important to stress that even with heterogeneity, the intuition remains the same. Aiding identification is the staggered introduction of consoles in the dataset: the PS2 was released a year earlier than the other two systems. Thus, whereas the change in sensitivity to price or software availability on the PS2 in the first year identifies the degree of consumer heterogeneity, whether or not these same early adopters are still "on the market" to subsequently purchase the Xbox or GC is a function of D.

4 Demand Estimation Results

Parameter estimates from the demand system are presented in Table 3. Multiple specifications are provided: column (i) estimates a static model without consumer heterogeneity (i.e., a standard logit model); (ii) and (iii) introduces dynamics without heterogeneity, but differ in whether or not consumers can singlehome or multihome; and finally, (iv) and (v) introduces consumer heterogeneity and dynamics when, again, consumers may singlehome or multihome. Estimation of models (i)–(iii) without any consumer heterogeneity is equivalent to estimating the hardware and software side sequentially in two separate stages; the nested fixed point routine introduced in the previous section to handle the selection of consumer heterogeneity is unnecessary. I will first describe demand results under the full model (v) before comparing across specifications.

[64] Recall that the utility from purchasing a new hardware system contingent on already owning a different console does not include software (both existent and forthcoming) developed for both the new and previously owned systems.

Nonlinear Parameter Estimates

All non-linear parameters $\theta_1 = \{\rho^{hw}, \rho^{sw}, \tilde{\alpha}_0^{p,sw}, \sigma_{p,sw}, \sigma_\gamma, \sigma_\epsilon, \beta_\gamma, D\}$ except complementarity factor D are estimated to be significant, including both parameters governing the variance in consumer heterogeneity. Signs of coefficients are as expected, with utility decreasing from price and having purchased a previous console. Regarding heterogeneity in price sensitivity, recall $\sigma_{p,sw}$ is the coefficient on a consumer's annual household income and not the standard deviation of the distribution, which explains its small magnitude ($.19 \times 10^{-5}$).

Heterogeneity in α^γ – a consumer's taste for software and gaming – is substantial. The estimated value of $\sigma_\gamma = .77$ indicates that a consumer at the 80% percentile of the distribution sees a game as 4 times cheaper than a consumer at the 20% percentile of α^γ – i.e., a game selling at \$50 for the intense gamer is seen as costing \$200, holding all else equal, for the less interested gamer. Thus, it is unsurprising that most consumers at the lower end of the distribution of α^γ do not purchase a console let alone many games. Figure 3 illustrates the estimated composition of the installed base of consumers across console by quintile of the α^γ distribution. For the first two years of each console's existence, over *half* of the users are in the top quintile of the distribution of α^γ; it is not until the end of the console's life-cycle that consumers with lower valuations begin to make up a significant share of users.

The ratio of the scale of the individual specific idiosyncratic error between the software and hardware side given by σ_ϵ is close to 2. This implies that uncertainty or idiosyncracies in preferences over a software title is twice as large as the idiosyncracies over the non-software component of a hardware system. Although hardware is substantially a larger purchase decision in dollar value, it has little if any stand-alone value apart from its software; thus, any idiosyncratic utility a consumer derives from a hardware-software system is primarily contained in the consumer's preference for different software titles. As a result, finding $\sigma_\epsilon > 1$ is not surprising.

Finally, σ_ϵ is also used to provide the ratio between mean hardware and software price sensitivities ($\alpha_0^{p,hw} = \sigma_\epsilon \alpha_0^{p,sw}$). The implied value of $\alpha_0^{p,hw}$ is reported at the bottom of the nonlinear parameter estimates. For robustness, I also estimated $\alpha_0^{p,hw}$ and $\sigma_{p,hw}$ separately from software parameters, and could not reject the restriction imposed by the model. Although the standard errors on $\alpha^{p,hw}$ were larger when estimated separately (due to the lack of price variation in the hardware side), the proximity of this estimate to the value imposed by the restriction supports the link between the hardware and software demand systems.

Linear Parameter Estimates

The bottom portion of Table 3 reports linear hardware and software parameters $\theta_2 = \{\alpha^x, \alpha^w\}$. Software title fixed effects and seasonality effects are provided in the next table. The first immediate observation is the large difference in estimated fixed effect for the PS2 as compared to its rival platforms. Despite controlling explicitly for software, installed base, age, and seasonality effects, the PS2 hardware unit (net price) is estimated to be 3 times as valuable as its closest hardware competitor. Several factors included in the PS2's fixed effect that are absent from its competitors include the ability to access and play the previous generation PS1's existing library of over 1,000 games (most released prior to the introduction of PS2), the ability to play DVDs right out of the box,[65] and possibly other factors including its unique hardware specifications, design aesthetic, and brand loyalty.

Unsurprisingly, the age of a console and software title is estimated to negatively affect lifetime utility from purchase. With hardware, it may partially reflect the fewer periods remaining to enjoy the console before the next generation of video game systems are released (i.e., obsolescence), or merely some form of decay with respect to its perceived value, quality, or desirability; with software, the latter effects seem more likely to be the reason. On the other hand, the observed installed base of a product positively impacts utility for the Xbox and GC, but not the PS2. Recall that the coefficient on installed base reflects how lifetime utility is affected by the *observed* installed base; insofar that expectations of future installed base are correct and completely accounted for in each product's fixed effect, the coefficient on installed base should be 0. As a consequence, the negative coefficient on the PS2's installed base component may merely be a consequence that initial estimates for the PS2's eventual installed base were much higher than what was observed, and over time perceptions (and hence, the utility) adjusted downwards despite an increase in the observed installed base. Similarly, a positive coefficient for the Xbox and GC may indicate a true immediate increase in utility from more people coming onboard a system or using a software title, or may be a result from consumer's revising their expectation over a product's eventual installed base upwards over time.

Table 4 reports month fixed effects for hardware and software. Seasonality, as expected, is dramatic in determining when people purchase goods with highly positive and significant coefficients on holiday months in particular.

Table 5 presents an OLS regression of recovered software title fixed effects $\{\alpha_k^w\}_{\forall k}$ on dummy variables indicating whether or not the title was exclusive (and if so, if it was published by a platform provider or not), the platform it was released on, and the month it was released. Results show exclusive titles that are published by the platform provider (which are typically developed in-house) tend to exhibit higher quality than average. This is highly significant across both industry veterans – Sony and Nintendo – who have had each at least one generation of prior experience in software development, but is not significant for new entrant Microsoft. This may be evidence for integration as effort or quality enhancing for those firms with experience; on the other hand, the possibility that first-party titles are selected upon before being acquired or that first-party studios are simply higher quality game developers cannot be ruled out either. There is, however, a significantly negative coefficient across platforms for third-party exclusive titles – that is, titles that chose to be exclusive voluntarily. This is unsurprising: most third-party titles that were exclusive did so not because they were compensated via some unobserved contract, but because the potential gains from multihoming would have been outweighed by the costs of porting to more consoles.

[65] Although the GC could not play DVDs, the Xbox required users to purchase a separate accessory to enable playback.

Fit of Model

In the appendix, I evaluate the fit of the estimated dynamic demand model, in particular focusing on assumptions 3.2 and 3.5 which govern the evolution of hardware mean-utilities $\{\delta_{j,t}\}_{\forall j,t}$ and product unobservables ξ, η. I find the parameterization of the first-order Markov process $F(\cdot)$ given by (16) provides a reasonable approximation of consumer expectations, errors ν^{hw}, ν^{sw} statistically appear to be uncorrelated and independent, and the degree of multihoming by consumers predicted by the model is consistent with industry figures.

Alternative Specifications and Preliminary Counterfactuals

I now return to the alternate specifications listed in table 3. There are three dimensions along which the five specifications differ: dynamics, heterogeneity, and consumer multihoming. The static specification in (i) does not allow for any dynamic considerations, which include the persistence of unobservable characteristics, consumers leaving the market after purchase, and forward looking agents; unsurprisingly, it has the smallest likelihood and poorest fit. For the rest of the specifications, results in the table agree on the significance and relative magnitudes for most parameters, and are similar for most parameters in θ_1 and θ_2. At the same time, a standard likelihood ratio test rejects the hypothesis that there is no consumer heterogeneity across specifications, and that consumers can only buy one console. Where do the differences come into play? Tables 6 – 9 present different predictions that arise from alternative specifications of the model. Here, substantive differences emerge.

Table 6 reports own and cross-price semi-elasticities for platforms across three specifications. Each cell reports the percent change in market share of the platform located in the column due to a permanent 10% decrease in the price of the row-platform, where "Outside" indicates substitution to or from the outside good.[66] The full model predicts much larger differences within own market shares following a price drop compared to estimates from other specifications – i.e., a permanent 10% price drop in the GC results in a predicted increase of 25% in total consoles sold instead of 19% – yet predicts such a price drop has a smaller impact on outside consumers. This latter effect occurs because in static model without multihoming, most of the increase in platform market shares from a price drop is predicted to come from consumers substituting away from the outside good and from other consoles; the full model, however, predicts that many of the new consumers are not substituting away from non-purchase or another console, but rather they are in fact purchasing an *additional* console.

Table 7 provides software own-price semi-elasticities for a representative "hit" title on each platform. These titles are selected since they are popular and released early in the lifetime of each console when the selection by consumers across platforms is most severe. The most striking difference again occurs across specifications which differ on multihoming: without multihoming, a price drop results in a large increase in the number of consumers who purchase a title since it induces

[66]Since platforms are active for multiple periods, the price change is assumed to apply across the entire time period, and market shares are computed from installed base figures at the end of the sample period (October 2005).

more people to purchase the console; with multihoming, this effect is significantly reduced (by over a half in the case of the Xbox and GC title) since many of those who might have purchased a console in order to access the title (i.e., high valuation consumers) would already have been predicted to own multiple consoles.

Table 8 presents changes in hardware installed bases if these three representative titles were not available – i.e., this provides an idea of the elasticity of demand with respect to a hit title.[67] These elasticities can be extremely large: Microsoft's *Halo*, in the full model, is predicted to have resulted in a 9% increase in the number of Xbox consoles sold (a difference of over 1.2M units). The dynamic specifications nearly double the predicted impact of a hit title on hardware demand vis a vis the static specification. This is not surprising: a static specification underestimates the quality of a product since it attributes non-purchase to low quality and not to people waiting for the price to fall. The addition of heterogeneity also impacts results: in its absence, a majority of consumers who substitute away from a console due to the loss of a hit title become non-purchasers of any console; however, with heterogeneity, these consumers instead are more likely predicted to substitute to another console. This is also intuitive.[68] Capturing these dynamics and predicting where consumers substitute to and from is crucial in understanding how and why platforms compete for exclusive software.

In an unreported specification, I also examined what would have happened if the same hit titles for Xbox and GC had still been available on those platforms, but had also multihomed instead of being first-party exclusives. In other words, this is the benefit Xbox and GC expected to receive from exclusivity. The same differences across specifications was observed as before, with the static specification predicting less of an effect and lack of heterogeneity overestimating the impact on the outside market shares. Nonetheless, I find that Xbox and GC would have been actually better off *losing* the title outright than having the title multihome – e.g., had *Halo* multihomed, Xbox would have seen its predicted installed base fall by over 12% instead of 9%, mainly since PS2 would have captured even more consumers as a result.

Finally, in table 9, I explore a "naive" counterfactual environment in which all titles are forced to be available on all consoles, holding all other observed hardware and software prices and characteristics fixed.[69] In such an environment, the differences across specifications become the most stark. All models without heterogeneity predict that when all titles are available on all consoles, the vast majority of non-purchasers become video game purchasers; in a dynamic context without heterogeneity, the model predicts that nearly every household would then purchase a videogame console. This is highly unrealistic – there are a significant number of consumers and households who, no mat-

[67] I restrict attention only to losing the particular title, and not any sequels or titles in the same franchise.

[68] E.g., if all consumers were the same, then those who do not purchase a console when that console lost a title would be no different than those who elected not to purchase a console in the first place. However, with heterogeneity, those consumers who substitute away from purchasing a console because it lost the title are already more predisposed to gaming (since they would have purchased a console in the first place), and thus are much more likely to purchase another console instead.

[69] How to extrapolate prices and qualities for a title onto consoles it was not previously released on is discussed in the next section.

ter what the availability of software may be, will not purchase a console due to income constraints or extreme preferences. Introducing heterogeneity helps to correct for this out-of-sample prediction, and estimates that only approximately another 20-25M consumers would purchase consoles in the event of "forced compatibility." Simulation runs also indicates that a model only allowing for consumer singlehoming underestimates both the benefit the PS2 receives from forced compatibility and the harm borne by the Xbox. Regardless of the specification, the counterfactual regime does seem to indicate that the PS2 does significantly better with the addition of new titles; the Xbox does significantly worse; and the GC does slightly better.

Nonetheless, this last counterfactual is hardly realistic; "forced compatibility" is not the appropriate regime to analyze when considering the absence of vertical integration or exclusive dealing. The absence of these exclusive arrangements does not necessarily result in all titles multihoming across all consoles. In such a case, some titles still may voluntarily elect to be exclusive (or just support two consoles) since the costs of supporting more may be prohibitive. Also, it is unlikely that the characteristics of all products would remain the same. To address these concerns, I next focus on software provision and the decision of which platforms to support.

5 Hardware-Software Network Formation

For the consumer demand analysis, the set of software products on each platform has been conditioned on in the data. However, when institutional features of an industry change – as is the case when hardware platforms cannot vertically integrate into software provision or offer contracts contingent on exclusivity – it is unlikely that the contracting relationships between parties will remain the same. To account for this, in this section I develop a model of software "demand" for hardware.

As discussed, each software title needs to be specifically developed for a particular console in order to compatible with it. If a title has agreed to an exclusive contract or is an integrated first-party title, the decision of which platform to support has already been made; otherwise, a third-party title typically chooses the set of platforms that maximizes its expected profits. Each title weighs two competing forces when deciding which platforms to support: on one hand, developing and releasing a title for a platform provides access to that console's base of users, which in turn may yield greater sales; on the other hand, such development requires the outlay of (non-negligible) "porting" costs which may or may not be recouped through these additional sales.

The estimates recovered from the demand system allow the econometrician to predict how many copies a title would have sold had it joined any set of platforms (including ones it was not observed to have done so), and how many more or fewer hardware consoles would have been sold as a result. However, a software title does not know exactly how many copies it will sell for sure when it makes its decision: unlike a consumer who can decide immediately which hardware platform or software title to purchase, a software developer must make a decision of which consoles to support $6-12$ months in advance of release. The challenge for the econometrician, thus, is to use estimates from the consumer demand system to construct estimates of each software title's *expected* profits.

However, when evaluating the expected profitability of joining any set of platforms, a software title has expectations over not only its own quality and price upon release, but also the number of consumers that will be onboard each platform. This introduces a significant complexity: in a platform market, a piece of software released for one platform will presumably induce more consumers to join, which in turn may induce more titles to join, hence driving more consumer adoption, and so on. As a consequence, every software firm must account for how other agents will not only act, but also how they will react to the title's own actions when it computes its own expected profits.

Nonetheless, by leveraging similar assumptions used in the consumer demand analysis that (i) software titles compete in independent markets and (ii) software titles perceive a sufficient statistic for hardware demand is the set of mean-hardware utilities $\{\delta_{j,t}\}_{\forall j,t}$, this particular issue can be resolved. In particular, a software title is affected by the actions of other titles only if they affect installed base of each console. By assumption, this can only occur through changes in $\{\delta_{j,t}\}$. I assume that each title believes it can shift the level of $\delta_{j,t}$ by deciding to join platform j, but does not change its perception over the evolution process of $\delta_{j,t}$ from its new state as a result of its own action. However, such beliefs over the evolution of $\{\delta_{j,t}\}$ are sufficient for every title to internalize the future responses of other agents.

In the next subsection I formalize this logic and detail the construction and computation of each software title's expected profits. Since porting costs are unobserved, I focus on their estimation and recovery in 5.2.

With these estimates and a means to compute the expected profits for each software title in hand, I turn to characterizing software demand for platforms. In 5.3, I present a dynamic game whereby all software titles are allowed to freely choose which platforms to develop for, and define and describe how to compute an associated equilibrium for the game. Finally, in 5.4, I test the fit of this model to the data by fixing the actions of all first-party titles, but allowing each third-party title to re-optimize and choose a new set of platforms. I find that the outcomes predicted by the model are very close to those realized in the data.

5.1 Software Expected Profits

Consider the decision faced by a particular software title k which will be released at time r_k. Assume τ months in advance, at time $r_k - \tau$, that title must choose a strategy $\mathbf{s}_k \in \mathbf{S} \equiv \{\{0,1\}^3\}$ which indicates which set of platforms k will develop for. For a given strategy $\mathbf{s}_k$, title k's expected discounted profits are given by (where, abusing notation slightly, $\mathbf{s}_k$ also represents the corresponding subset of J):

$$E[\pi_k(\mathbf{s}_k;\theta_C)|\Omega_{k,r_k-\tau}] = E[\Bigg(\sum_{t=r_k}^{T} \beta^{\tau+t-r_k} \sum_{j\in\mathbf{s}_k} Q_{j,k,t}((1-rmkup)p_{j,k,t} - mc_j)\Bigg)|\Omega_{k,r_k-\tau}] - C_k(\mathbf{s}_k;\theta_C) \tag{18}$$

where $Q_{j,k,t}$ is the quantity of title k sold on platform j at time t, *rmkup* denotes the markup captured by retailers, mc_j is the marginal cost of production on console j (which includes royalties paid to the platform provider), and $C_k(\mathbf{s}_k;\theta_C)$ are the "porting" costs of producing title k for all the platforms within $\mathbf{s}_k$ which depends on some vector of parameters θ_C. In addition to development and programming costs, $C_k(\cdot)$ contains all other fixed costs related to the production of the game including distribution and marketing. These porting costs are known to each software title but not to the econometrician. Finally, expectations are conditional on $\Omega_{k,r_k-\tau}$, software title k's information set at time $r_k-\tau$, which includes any factors affecting market characteristics and consumer demand.

From the demand system, $Q_{j,k,t}$ can be computed for any title conditional on knowing the installed base of consumers onboard platform j at time t (who have not yet purchased k), and the quality and price of that software title. I will assume firms share the same beliefs as consumers over the evolution of both $\{\delta_{j,t}\}_{\forall j,t>r_k-\tau}$ and $\{\zeta_{j,k,t}\}_{\forall j,t>r_k-\tau}$, given by processes $F(\cdot)$ and $\{G_j(\cdot)\}_{\forall j}$ respectively. In turn, the expectations will be consistent with the empirical distribution of these parameters. Essentially, this implies that a software title does not explicitly consider the strategy choices of other titles; rather, it only does so only through how they enter and affect the mean-utilities of each hardware console, δ. As long as title k knows the transition probabilities $F(\cdot)$ and $\{G_j(\cdot)\}_{\forall j}$, then the installed base on each platform and the mean-utility of each hardware platform at time $r_k-\tau$, as well as the title's own starting quality are all that are required to compute the expected number of titles $Q_{j,k,t}$ sold onboard each platform. In the appendix, I further detail the assumptions and associated computational routine required to recover these estimates.

Note that for the purposes of this analysis, platforms are not assumed to be strategic agents other than setting the prices for their own consoles (which is internalized in the evolution of $\delta_{j,t}$). This is motivated by the desire to analyze an environment without exclusive deals or vertical integration; as such, I rule out any preferential treatment by platform providers since these deals are often made in exchange for exclusivity. Furthermore, platforms typically pre-announce and commit to royalty rates charged to third-party software developers in advance of a system's release.[70]

Finally, for estimation purposes, I will assume the retail cost margin is fixed at 35% and marginal costs are constant across platforms at \$10 (reflecting royalty rates of approximately \$7 and production costs of \$3 per game disc). These figures are consistent with information provided by industry and public sources.[71]

5.2 Recovery of Development and Porting Costs

In order to compute a title's expected profits from choosing any particular action $\mathbf{s}_k$, one final issue remains: development and "porting" costs $C_k(\cdot;\theta_C)$ are unobserved. To estimate and recover these unobserved costs, I proceed using techniques developed in Pakes, Porter, Ho, and Ishii (2006) which

[70]See e.g. Kent (2001), Hagiu (2006).
[71]See e.g. Takahasi (2002).

uses a methods of moments estimator based on inequality constraints.[72]

Consider again the decision of a third-party title k who decided τ months in advance of release which platforms to develop for. The key assumption used to generate the moments for estimation is that for each title observed in the data that was brought to market by a third-party publisher, the expected profits from developing for the set of platforms it did should have been higher than developing for any other set *holding the actions fixed for all titles released up to that point in time*:

Assumption 5.1. *For each third-party software title k, the observed choice of platforms $\mathbf{s}_k^o$ maximized its expected profits:*

$$E[\pi_k(\mathbf{s}_k^o;\theta_C)|\Omega_{k,r_k-\tau}^o] \geq E[\pi_k(\mathbf{s}_k';\theta_C)|\Omega_{k,r_k-\tau}^o] \;\forall\; \mathbf{s}_k' \in \mathbf{S}$$

where $\Omega_{k,r_k-\tau}^o$ denotes the *observed* state of each title's information set at time $r_k - \tau$.[73] Also note that I am implicitly assuming that the choice for any title released by a third-party publisher is made independently of any other title.

I use a parsimonious specification for $C_k(s_k;\theta_C)$:

$$C_k(\mathbf{s}_k;\theta_C) = c_0(\mathbf{s}_k) + \sum_{j\in\mathbf{s}_k} c_j\alpha_{0,j,k}^w + \nu_k^c \tag{19}$$

where $\alpha_{0,j,k}^w$ represents the mean software fixed effect for title k on platform j, estimated from the demand side, and $\theta_C \equiv \{\{c_0(\mathbf{s})\}_{\forall\mathbf{s}\in\mathbf{S}\setminus\{0\}^3}, \{c_j\}_{\forall j}\}$.[74] ν_k^c represents title-specific costs that affect all strategy choices equally. The difference in costs between two different titles are thus assumed to be contained within differences in the estimated software fixed effect and some unobservable title-specific component.

Given assumption 5.1, the expected difference in profits across all software titles between the observed strategy chosen and any alternative should be positive:

$$E_k[E[\pi_k(\mathbf{s}_k^o;\theta_C)|\Omega_{k,r_k-\tau}] - E[\pi_k(\mathbf{s}';\theta_C)|\Omega_{k,r_k-\tau}]] \geq 0 \;\forall\; \mathbf{s}' \in \{\mathbf{S}\setminus\{0,0,0\}\}$$

Since I do not observe software products which are not released on any platform, I restrict attention to strategies that involve joining at least one platform.

Let $\mathbf{K}_\mathbf{s}$ denote the set of titles that choose strategy $\mathbf{s}$. For each $\mathbf{s} \neq \{0\}^3$ and $\mathbf{s}' \notin \{\mathbf{s}\cup\{0\}^3\}$,

[72]See Ho (2007) and Ishii (2005) for other applications of this framework. Unlike using a multinomial logit model to determine costs based on observed actions for each software title, this procedure can handle the correlation of certain kinds of measurement and expectational errors without needing to fully specify their joint distribution. Furthermore, it imposes fewer assumptions than would otherwise be required.

[73]I will assume that the econometrician's estimate and a title's estimate of expected profits are the same. As long as the error between the two is mean zero across titles and strategy choices and independent of any instruments chosen, the following analysis does not change (this error is equivalent to the ν_1 error discussed in Pakes, Porter, Ho, and Ishii (2006)).

[74]In estimation, values of $\alpha_{0,j,k}^w$ are scaled to be within $[1, 10]$.

converting expectations into sample means yields the following inequality moments:

$$\frac{\sqrt{\#\mathbf{K_s}}}{\#\mathbf{K_s}} \sum_{k \in \mathbf{K_s}} (E^o \pi_k(\mathbf{s}; \theta_C) - E^o \pi_k(\mathbf{s}'; \theta_C)) \otimes g(\omega_{k,t-\tau}) \geq 0 \tag{20}$$

for any $\omega_{k,r_k-\tau} \in \Omega_{k,r_k-\tau}$, where $\otimes$ represents the Kronecker product and $g(\cdot)$ is any positive valued function. I weight by the square root of the number titles that choose each particular strategy $\mathbf{s}$ in order to account for the fact that there should be less expectational noise in expected profits for strategies chosen by many titles.

For now, equation (20) defines 42 inequalities (7 non-zero strategies, each with 6 alternative strategy comparisons) to be used in estimation. If there are multiple values of θ_C which satisfy these inequalities, all are admissible and a set estimate is obtained; otherwise, the value $\hat{\theta_C}$ which minimizes the absolute value of deviations in the inequalities will be used.[75]

Since only strategies that involve joining at least one platform are compared, all components in θ_C are not identified: only the relative differences between $c_0(\mathbf{s})$ and $c_0(\mathbf{s}')$ can be determined. Nonetheless, for the purposes of the subsequent analysis, relative differences are all that are required in order to determine the optimal choices for software titles. In estimation, $c_0(\{1,0,0\})$ (i.e., the constant cost for developing only for the PS2) is fixed to be 0.

Estimates

Table 11 presents porting cost estimates for $\tau = 6$ (i.e., software titles make their decision 6 months prior to release). Since the costs for developing solely for the PS2 are fixed to be 0, these estimates reflect the *relative* costs of porting to a particular set of consoles.

The specification reported allows costs to vary across different software genres. This is important because certain consoles may be more difficult to develop for than others in particular genres; similarly, higher quality games also may be more or less difficult to produce depending on the genre. The results confirm such variation.

For the average title in all genres, magnitudes are in line with what is reasonable – developing for two consoles is generally more expensive than developing for one, but still cheaper than developing for all three; and developing for the Xbox and GC is cheaper than doing so for either the PS2 and GC or the Xbox and GC. Furthermore, results predict that for most titles, Xbox and GC are to be significantly cheaper to develop for than the PS2 – this is also consistent with institutional details.[76]

[75] When constructing the inequality estimators, I also omit "high-quality" exclusive titles brought to market by third-party publishers, which I assume to be those with estimated fixed-effects $\alpha^w_{0,k} > -3$. The reason for this restriction is that these exclusive titles, although not first-party, may have been subject to unobserved exclusive deals involving royalty rate reductions, lump sum payments, development assistance, or joint marketing promotions. The underlying assumption is that all other titles – those that multihomed were or of low enough quality – did not receive any exclusive contracts or preferential treatment from console providers. Because of the specification of costs in (19), any title specific errors are differenced out and no selection problem is introduced.

[76] E.g., the PS2 with a new CPU architecture had a reputation of being difficult to develop for, whereas the Xbox was essentially an Windows-Intel PC using APIs many developers were already familiar with.

Repeating the exercise for different values of τ (including 9 and 12 months) did not significantly affect estimates of porting costs, nor affected the counterfactual analysis that follows. The choice of instruments, however, did affect point estimates and associated confidence intervals. To account for this variation, I used several different sets of instruments and associated porting cost estimates when computing confidence intervals for all of the following counterfactual exercises. As will be evident, the predictions and results were still fairly robust to the choice of instruments used.

5.3 Dynamic Network Formation Game

I now present a dynamic network formation game whereby each software product prior to release selects which platforms to develop for having formed expectations over the future profitability of each potential strategy. The setup allows for contracting partners and consumer demand to change over time with past actions influencing future decisions, a crucial feature to capture in this networked industry.

Of course, other industry features may change as well, including the investment incentives which will affect entry and exit of new hardware and software products. For the purposes of expositing the model, I will here focus only on the changes in contracting partners, assuming that the set of available software products is given and porting costs, royalty rates, and retailer markups do not change. In the next section when I estimate a counterfactual regime, I discuss potential ways of relaxing some of these restrictions.

Setup and Timing

In each period t, assume there is a set of software products $\mathbf{K}_{R(t)}$ that will be released on at least one console. At time $t-\tau$, every title $k \in \mathbf{K}_{R(t)}$ simultaneously commits to the set of consoles that it will be released on τ periods in the future. This decision is private, and will not be observable or known to any other industry participant until t. Furthermore, software participants do not have any information about future software releases or availability $\mathbf{K}_{R(t')}\ \forall\, t' > t-\tau$. Finally, platforms have no strategic actions, and royalty rates are fixed and common knowledge before any title makes a decision.

To be clear, at each period t the timing of actions is assumed to be as follows:

1. All titles $k \in \mathbf{K}_{R(t)}$ are released and added to the stock of existing software products;
2. All products' characteristics (which are subsumed by $\delta_{j,t}$, $\zeta_{k,t}$) are determined;
3. Consumers make hardware and software adoption decisions;
4. All titles $k \in \mathbf{K}_{R(t)+\tau}$ choose which platforms to join.

I will assume title's expected discounted profits from joining platforms $j \in \mathbf{s}_k$ is given by (18).

Equilibrium and Computation

An equilibrium of this game will be a set of strategies $\{\hat{\mathbf{s}}_k\}_{\forall k}$ and a set of Markov transition processes F and $\{G_j\}_{\forall j}$ over the evolution of $\{\delta_{j,t}\}_{\forall j,t}$ and $\{\zeta_{j,k,t}\}_{\forall j,k,t}$ such that:

1. For every k, $\hat{s}_k$ maximizes its expected profits time $r_k - \tau$:

$$E[\pi_k(\hat{\mathbf{s}}_k;\theta_C)|\Omega_{k,r_k-\tau}] \geq E[\pi_k(\mathbf{s}'_k;\theta_C)|\Omega_{k,r_k-\tau}] \; \forall \; \mathbf{s}'_k \in \mathbf{S} \setminus \{0,0,0\}$$

2. Transition processes F and $\{G_j\}_{\forall j}$ are consistent with realized values $\{\delta_{j,t}\}_{\forall j,t}$ and $\{\zeta_{j,k,t}\}_{\forall j,k,t}$ implied by $\{\hat{\mathbf{s}}_k\}_{\forall k}$.

A few comments are in order. First, F and $\{G_j\}_{\forall j}$ will likely be different from the one estimated in the demand system since a rational expectations equilibrium accounts for any changes from the original set of contracting partners. Secondly, recall a software title does not explicitly consider the strategy choices of other titles; rather, it only does so only through how they enter and affect the mean-utilities of each hardware console, δ. In this regards, the equilibrium has certain similarities with the notion of *oblivious equilibrium* developed by Weintraub, Benkard, and Van Roy (2007) in Ericson and Pakes (1995)-style models, whereby agents only condition on its own state and knowledge of the long run average industry state.

Finally, it may be the case that there are multiple equilibria: different beliefs over the evolution of each hardware platform's quality may sustain different actions which in turn rationalize those beliefs. Although in computation I restrict attention to space of beliefs given by the parameterizations of F and G_j used on the demand side (given by equations (16) and (17)), there still may be multiple equilibria which satisfy the conditions given above. However, there are bounds on the utility a consumer obtains from each platform – i.e., there is are minimum and maximum values $\underline{\delta}_j$ and $\overline{\delta}_j$ for each platform which correspond to hardware mean-utilities without any software titles or with all software titles onboard – and thus there are significant restrictions on the set of sustainable beliefs in any equilibrium. To account for the possibility that there may still be multiple equilibria, I run the following computation algorithm with different starting beliefs:

- I first fix the transition processes governing the evolution of δ and ζ to starting beliefs F^0 and $\{G^0_j\}_{\forall j}$. For robustness, I use 5 different sets of starting beliefs F^0 which govern the evolution of hardware qualities δ: one which assumes no software title joins any console, one which all titles join every console, and three different sets in which all titles join only one console.

- In each iteration n, I proceed from $t = 0$ forward and at every period: update $\delta^n_{j,t}$ for each console based on the set of new titles released; evaluate consumer demand over the set of hardware and software products; and compute the optimal strategy s_k for each title $k \in \mathbf{K}_{t-\tau}$.[77] Once all periods and titles have selected a strategy, I use the implied paths of

[77] Despite $\mathbf{s}^*$ being defined as a correspondence, in computation there is always a optimal strategy for every title. The appendix provides details on computing expected profits.

$\{\delta^n_{j,t}\}$ and $\{\zeta^n_{j,k,t}\}$ to update the transition processes according to the regressions (16) and (17), obtain new estimates for F^{n+1} and $\{G^{n+1}_j\}_{\forall j}$, and repeat until no software title changes its chosen action from the previous iteration and the estimated transition processes converge.

5.4 Fit of Dynamic Network Formation Model

To evaluate how well this model predicts the strategic actions of firms, I estimate an equilibrium holding fixed the actions of all first-party software titles, and allow each third-party title to re-optimize and choose a new set of platforms to support.

Table 12 presents a comparison between the observed data and those predicted by the new equilibrium model. Confidence intervals are constructed by repeating the simulation using draws from the distribution of estimated porting costs in the previous section. The model predicts both installed base figures and market shares for each console to be very close to those observed in the data.

Although the PS2 is predicted to have fewer titles join and the Xbox and GC more, restricting attention to only "hit" titles – defined as titles selling over 100K or 1M copies on a given console – indicates a much closer fit. This indicates that although the estimation error or specification error in porting costs may likely affect the actions chosen by small titles, it becomes less of an issue for the larger, more popular games. Since the actions of these larger titles are the only ones that dramatically affect platform market shares, as long as their actions are accurately predicted, the estimated aggregate industry figures such as market shares and installed base figures will be similar to those observed.

Finally, the total number of titles sold on each platform is also provided, and again the model predicts figures close to those given by the data. These totals are important since most of each platform provider's profits are derived from royalty payments on software, not from hardware sales; they translate directly into the profitability and success of each platform.

In all simulated runs, using different starting beliefs did not change the computed equilibrium. This is not surprising: note that the decision of which consoles to support is typically robust to small fluctuations in beliefs over the evolution of δ. Only "high quality" titles can shift the value of δ for any given console in any significant way, and they will usually choose their strategy regardless of what they believe other software titles will do; most of a title's sales occur in the first 3 months of release, and as long as there are sufficient numbers of consumers on board each platform at that given moment, these titles will typically multihome. On the other hand, for most mid and low-quality titles, the first-order impact of porting costs on profits typically dwarfs the impact of any small shifts in installed bases across consoles implied by changes in δ.

6 Policy Experiment: Banning Vertical Integration and Exclusive Contracting

This section uses the estimates obtained so far to compute a counterfactual environment whereby console providers are prevented from integrating into software development and producing first party titles, and are unable to offer exclusive contracts to third-party titles. However, third-party titles still may voluntarily choose to be exclusive if they find the costs of porting too high.

As mentioned before, such a change in the contracting space may also have affected the entry or exit of software products. For example, integration may have encouraged investment in first-party software that otherwise would not have been produced. To account for the possibility that there might be changes in the space of software products, I consider two cases: (i) first-party titles are assumed to still enter the market as third-party titles; (ii) all first party titles no longer are produced and thus do not enter. These two alternatives can be used as potential proxies for what actually may occur absent exclusivity and integration.

Some caveats still remain. Although software providers form expectations over the future quality and prices of products when deciding which platforms to join, I assume prices are the same as observed price paths in the data when computing the outcomes of the counterfactual regimes. Without a full model governing the dynamic price setting behavior of firms, these counterfactual price paths become difficult to determine.[78]

6.1 Industry Structure

The results from the counterfactual simulations are presented in table 13. Generally, more titles are predicted to multihome and port than observed in the data, which is expected. In the first specification when former first-party titles still enter the market as third-party titles, the change in market shares is stark: Sony's PS2 is predicted to command over 75% of the market by October 2005 as opposed to half, and sells nearly double the number of consoles and nearly 5 times as many total copies of software as before. On the other hand, the Xbox does significantly worse, selling 5M fewer consoles and nearly half as many copies of games. The relative success of Nintendo's GC is less clear; it loses market share and sells approximately the same number of hardware devices, but does manage to sell more software titles. The reason that the response of the Xbox and GC are different is because the Xbox, which sells at the same price point as the PS2, is more of a direct competitor and substitute for the PS2. Without exclusive titles distinguishing the PS2 from the Xbox, most consumers will select the PS2 due to its higher predicted fixed effect and 14 month head start in accumulating a larger installed base and software library. However, consumers may still flock to the GC because of its cheaper price – it appeals to price sensitive consumers as a first

[78] Nair (2007) provides one possible approach to modelling a firm's dynamic pricing decision. Nesting this type of optimization problem within this paper's dynamic hardware and software demand system while still accounting for the endogenous selection of heterogenous consumers across platforms is too computationally burdensome for implementation. Nonetheless, I also compute the counterfactual simulations assuming software prices follow a first-order Markov process estimated in the appendix, and find that results do not change significantly.

console, and also to other owners of the PS2 or Xbox for its titles that may still be voluntarily exclusive or potentially of higher quality.[79]

The second specification conducts the same exercise, except now removes all former first-party titles from consideration. The industry as a whole does worse off since most of the first-party titles were blockbuster hits; nonetheless, because the PS2 still captures a majority of the market and enough software products, it still does better than it did originally. Here, both the Xbox and GC are predicted to do significantly worse, with each selling far fewer consoles and software titles as they did in the presence of exclusivity. Recall table 2 showed nine of the top ten titles sold on GC were first party titles, and that table 5 indicated the first party titles on GC were shown to be much higher quality on average; if the inability to vertically integrate leads to the absence of these former first party titles, then it is unsurprising that the total number of titles sold on the GC is predicted to drop precipitously in their absence.

6.2 Consumer Welfare

By removing exclusive arrangements, users onboard each platform are predicted to have access to a greater selection of high quality titles. To analyze consumer welfare gains, I calculate the compensating variation for consumers who are predicted to purchase a console in each counterfactual environment.

In the first regime where first-party titles still enter the market, I find that total consumer welfare increases by approximately $7B *holding fixed hardware and software prices and entry.* Half of this increase is realized by those consumers who would have purchased a videogame console in the previous regime; these consumers receive on average approximately $72 more in surplus. The other half of consumer welfare increase is realized by new consumers who previously did not purchase a console, but now would; 26M new households on average receive approximately $140 in surplus. However, in the second regime whereby first-party titles are no longer produced, consumer welfare is predicted to increase only by $.7B with existing purchasers gaining on average only $10, and new purchasers (of which there are far fewer) approximately $80. The loss of first-party titles, thus, drastically reduces the potential gain from increased software compatibility.

Both of these calculations again ignore the possibility that Microsoft or Nintendo may have exited the market, or Sony, with its increased market power, may have increased prices. For example, in the first regime if Microsoft and Nintendo did not enter at all due to the inability to integrate or obtain exclusive titles, total consumer welfare would be predicted to have *fallen* by $1B despite having all titles onboard a single console. Similarly, I find Sony – holding the prices fixed for the Xbox and GC – would have found it profitable to raise the price of its PS2 unit by over $200, again eliminating most consumer welfare gains in either counterfactual regime.

[79] For titles released on multiple consoles on the data, different platform-title specific fixed effects are estimated. As a result, the utility from the same title may be different onboard different consoles.

6.3 Discussion and Policy Implications

Counterfactual experiments suggest that vertical integration and exclusive contracts were generally pro-competitive at the platform level, benefiting Microsoft and Nintendo in terms of market share and helping them establish a substantial foothold into the sixth-generation videogame market. The PS2 had already captured an installed base of 5M users before its two competitors entered a year later; as a result, without exclusivity, the Xbox and GC would likely only have been able to induce a developer to release a title for their respective console after a version had already been developed for the PS2.

Furthermore, with Sony commanding over three-quarters of the market, it is likely that it could have sustained higher prices – although it may not have necessarily charged a higher entry price, it would nevertheless not have been as hard pressed to anticipate or match Microsoft's and Nintendo's console price cuts in 2001 and 2003. Combine this with the possibility that Microsoft and Nintendo may not have found it profitable to remain in the videogame industry and either exited or not produced their seventh generation consoles, any immediate consumer gains from increased access to software may very likely have been offset by these dynamic consequences of monopolization. Thus, although such welfare calculations are sensitive to the assumptions used concerning entry, exit, and price setting behavior, the implications governing industry structure, market concentration, and platform competition appear to be robust.

It is worth stressing that although exclusive arrangements may have encouraged *platform* competition, this does not necessarily imply that they encouraged *software* competition. In the videogame industry, without modelling the entry and exit of new titles, the effect on software is ambiguous – having only a single monopoly platform to support might have reduced porting costs required for a third-party developer since only one console would have to be developed for, but a more competitive environment with multiple integrated platform providers might have increased investment in first-party software or led to fiercer competition among platforms for titles through reduced royalty rates or increased development and marketing assistance.

In other industries, the impact of vertical integration and exclusivity on software competition is more clear. For example, consider Microsoft's integration of its "hardware" (Windows OS) into the browser and media application space. The courts in both *US v. Microsoft* and *European Union v. Microsoft* ruled that these actions resulted in foreclosure of competing software vendors (e.g., Netscape and Real Networks). Although both the videogame industry and the PC industry are similar hardware-software environments, the fact that PC applications are very close substitutes (consumers typically only use one word processor, browser, media player, or spreadsheet program), whereas videogames are not (buying one action game does not preclude the purchase of another), indicates that "upstream" software foreclosure may be more of a concern when such substitutability is an issue. The results of this analysis, if extended to the PC industry, would likely indicate that vertical integration and exclusivity into software development may aid other *platform providers* (such as MacOS and Linux) in competing against Windows.

A more suitable comparison to the videogame industry is in television distribution, specifically with regards to competition between satellite and cable providers for exclusive content. In the US, DirecTV's exclusive contracts with certain content providers – notably with the National Football League for a package of its out-of-market games – substantially contributed to its success and ability to induce consumers to substitute away from cable. The impact of this competition was substantial: Goolsbee and Petrin (2004) estimate that entry by satellite providers reduced cable prices by about 15% and encouraged improvements in cable quality, yielding aggregate welfare gains of approximately $5B. In this regards, the U.S. Senate's recent actions which prevented an exclusive deal between Major League Baseball and DirecTV – as detailed in footnote 9, an intervention motivated mainly by a static efficiency desire to expand consumer access – may have negatively affected competition in the industry. Without these exclusive carriage deals, cable providers as well as other incumbent distribution channels in similar media markets would have faced less competitive pressure.

At the same time, in other industries platform competition may not be desirable and monopolization actually preferable. Often this is true when a dominant platform provider cannot raise prices or otherwise exercise market power. For example, consider the two main competing standards for next generation DVDs – Blu-ray and HD-DVD, sponsored primarily by Sony and Toshiba, respectively. Here, even if one standard wins and "monopolizes" the market, that platform sponsor cannot increase prices since it has already committed to licenses and royalty rates with hardware manufacturers and movie studios. At the same time, having a single clear standard emerge as the dominant platform will effectively remove uncertainty from the marketplace and likely spur consumer adoption, thereby increasing total welfare.[80] As a consequence, integration and exclusive contracting between the standard sponsors and motion picture studios (e.g., Sony's ownership of Columbia Pictures or Toshiba's exclusive marketing deal with Paramount) – although, according to the analysis of this paper, "pro-competitive" in encouraging the existence of multiple formats – may actually be contributing to an undesirable and lengthy standards battle.

7 Concluding Remarks

This paper has shown that integration and exclusive contracting between console manufacturers and software developers in the videogame industry likely encouraged and enabled platform competition. Evidence suggests that in this and other platform markets where upstream foreclosure is not a concern and where forced exclusivity contracts are not permitted, antitrust intervention or regulation may not be necessary. Furthermore, when evaluating the possibility of foreclosure or entry deterrence in dynamic networked environments, traditional static analysis may fail to uncover significant pro-competitive effects of exclusive vertical arrangements.

This paper focused on developing the methodological tools and framework to empirically measure the impact of these exclusive arrangements. On the consumer side, I developed a demand system that could handle the dynamic selection of forward-looking, heterogeneity agents across

[80]See, e.g., Farrell and Saloner (1985) and Katz and Shapiro (1986).

platforms. On the software supply side, I detailed and estimated a computationally tractable dynamic network formation game which allowed agent's to anticipate the future actions of other players by conditioning on them through a small dimensional set of state variables.

Although counterfactual simulations indicate that the industry is far more competitive when exclusive arrangements are allowed, it is true that consumers may immediately benefit from the access to a wider selection of software titles onboard a platform when these arrangements are prohibited. Whether or not consumers would have been better or worse off as a whole depends crucially on whether the incumbent would have raised prices in response to its monopoly power, and whether or not competing platform providers would have exited. Either would be sufficient to eliminate any potential consumer welfare gains.

To make analysis tractable, a few key assumptions were made and carried throughout the paper. First, I assumed the independence of software titles; in other types of platform industries, there may be stronger substitution effects across upstream products. Secondly, to model the dynamic decisions for both consumers and firms, I relied on assuming product characteristics evolved according to Markov processes consistent with their realized empirical distributions. Finally, I abstracted away from software entry and exit decisions as well as the (dynamic) platform pricing problem. Some robustness checks and alternative formulations were explored, and it is unlikely that relaxing these assumptions would dramatically change the results of the counterfactual exercises. However, completely addressing these issues is beyond the scope of this paper and the subject of future work.

A Demand System: Further Details

This section of the appendix provides further details on the specification, estimation, computation, and robustness checks for the consumer demand analysis.

A.1 Multiple Hardware Purchases

To reduce the number of options a consumer has at any period in time, I will assume that a consumer can purchase at most one console per period, will never purchase a console she has already bought, and can return in any future period to purchase another console. The consumer's value function for being able to purchase a new console is thus given by:

$$\begin{aligned} V_i(I_{i,t}, \epsilon_{i,t}, \Omega_{i,t}) \quad = \quad & \max\{\max_{j \notin I_{i,t}} u_{ijt}(I_{i,t}, \omega_{i,t}(\Omega_{i,t})) + \beta E[V_i(I_{i,t} \cup \{j\}, \epsilon_{i,t+1}, \Omega_{i,t+1}|\Omega_{i,t}], \\ & u_{i,0,t}(I_{i,t}) + \beta E[V_i(I_{i,t}, \epsilon_{i,t+1}, \Omega_{i,t+1})|\Omega_{i,t}]\} \end{aligned} \tag{21}$$

That is, a consumer will either choose to purchase or not purchase a new console, and if she does decide to buy, she will purchase the one that delivers the highest expected lifetime utility. In either case, she accounts for the continuation or option value of remaining on the market and updates her inventory state depending on her chosen action.

I modify assumption 3.2 to account for the dependence of mean-hardware utility on inventory:

Assumption A.1. *For any inventory state $I_{i,t}$, (consumers perceive that) $\{\delta_{ijt}(I_{i,t})\}_{j\in J_t}$ can be summarized by an exogenous first-order Markov Process:*

$$F(\{\delta_{i,j,t+1}(I_{i,t})\}_{j\in J_{t+1}}|\Omega_{i,t}) = F_i(\{\delta_{i,j,t+1}(I_{i,t})\}_{j\in J_{t+1}}|\{\delta_{i,j,t}(I_{i,t})\}_{j\in J_t}, I_{i,t}, m(t)) \tag{22}$$

where now the evolution is not only individual specific, but inventory state specific as well. Using this assumption and assumption 3.1 on ϵ, (21) can be integrated over ϵ and rewritten as:

$$\begin{aligned} & EV_i(\{\delta_{i,j,t}(I_{i,t})\}_{j\in J_t}, I_{i,t}, m(t)) = \\ & \ln\Bigg(\sum_{j'\notin I_{i,t}} \big(\exp(\delta_{i,j',t}(I_{i,t}) + EV_i(\{\delta_{i,j,t}(I_{i,t}\cup\{j'\})\}_{j\in J_t}, I_{i,t}\cup\{j'\}, m(t+1)))\big) \\ & \qquad + \exp(\beta E[EV_i(\{\delta_{i,j,t+1}(I_{i,t})\}_{j\in J_t}, m(t+1))|\{\delta_{i,j,t}(I_{i,t})\}_{j\in J_t}, m(t)])\Bigg) \end{aligned} \tag{23}$$

The only remaining issue is that the equations which govern the predicted share of consumers which purchase a console must now integrate over the distribution of not only consumer types, but consumer inventory states as well. Accounting for this as well as the evolution of consumer inventories over time are provided in A.3 when computational details are discussed.

A final concern governs the software purchase decision of consumers who own multiple consoles – a consumer who has already purchased a title on one console she owns should not typically also purchase the same title on a second console. Unfortunately, correcting for this would involve tracking the inventory of software purchases for each consumer, which is not computationally feasible. Instead, I will assume that if a title is released on multiple consoles, than any consumer who owns multiple consoles is equally as likely to purchase the title on any particular console, and adjust the potential market size each title faces accordingly.[81]

A.2 Other Industry-Specific Issues

There are a few remaining institutional specifics that affect the estimation of the model.

[81] E.g., consider a title released on both the PS2 and Xbox. For demand faced on the PS2, the title is assumed to face a market of all consumers who only own the PS2 plus one-half of those consumers who own both the PS2 and Xbox.

The first involves an issue of staggered platform release dates: Sony's PS2 console was released in October 2000, whereas the Nintendo GC and Microsoft Xbox were not released until November 2001. Consequently, for a portion of the data, only one console was available. Nonetheless, consumers knew that the GC and Xbox would be released during the 2001 holiday season even a year in advance of the actual release.[82] As a result, I model the consumer's relevant problem from October 2000 to October 2001 as a finite horizon optimal stopping problem with only one hardware console available, and assume that consumers know the starting values for the new products when they are introduced in November 2001.

The second is that the PS2 is backwards compatible with titles released for Sony's previous generation console, the original Playstation (PS1). Any utility derived from titles released for the PS1 prior to October 2000 as well as expectations over future software availability would be subsumed in the PS2's fixed-effect; however, any unexpected utility from PS1 titles released afterwards would not be accounted for. From the release of the PS2 in October 2000, there were 387 titles released for the PS1, 332 of which were not also released for the PS2. None of these were large successes. Since it is impossible to differentiate whether or not purchasers of these software titles owned a PS1, PS2, or perhaps even both, I will assume that these titles do not influence a consumer's decision to purchase a PS2. It seems reasonable that a consumer interested in playing older generation titles either would already own one or purchase the much cheaper console, and that the role of these new PS1 titles on the decision to purchase a PS2 is marginal at best.

Thirdly, the PS2 exhibited shortages during the first few months of its launch and supply was not able to meet demand, and without correction the model would potentially predict a lower value for $\delta_{i,j,t}$ than the true value for those early months. However, if I assume that during these months access to the console was independent of consumer heterogeneity (and consumers purchased in the same proportion had there not been a shortage), then ignoring the implied $\delta_{i,j,t}$ for the first few months during estimation would still yield consistent results.[83] Doing so did not significantly affect results.

A.3 Computational Details

For notational purposes, let ι index the "inventory-state," where $\iota \in \{0,1\}^3$ and denotes which platforms are owned in each state. Slightly abusing notation, $\iota = 0$ will indicate no platforms are owned; $\iota = 7$ that all 3 have been purchased. As noted before, I discretize the distributions of α^p and α^γ to model consumer heterogeneity. Consequently, consumers can be divided among R groups (indexed by i, each with price-sensitivity and gaming-preference coefficients $(\alpha_i^p, \alpha_i^\gamma)$ and initial population share $\lambda_{i,t=0,\iota=0}$ according to the distributional assumptions given. At each period, the fraction of each type of consumer on a given console is given by $\lambda_{i,t,\iota}^j$. For estimation, α^γ takes on 11 distinct values, and α^p has 5, resulting in $R = 55$ distinct consumer types.

As discussed, the estimation routine has the following main steps:

- To evaluate a candidate θ, iterate on the following until convergence on set of software utilities $\{\Gamma_{j,t}\}_{\forall j,t}$ is obtained:
  - i Hardware Adoption: For a given $\{\Gamma_{j,t}^n\}_{\forall j,t}$, determine mean console utilities $\{\delta_{i,j,t,\iota}\}_{\forall i,j,t,\iota}$ which match observed shares in data with those predicted by the model. Also obtain distribution of consumer types who have purchased a given console at any period of time.
  - ii Software Adoption: Given the distribution of consumers onboard any hardware platform, compute mean software utilities $\{\zeta_{j,k,t}\}_{\forall j,k,t}$ for every software title on every platform that, again, match observed shares in data with those predicted by the model. Update implied software utilities $\{\Gamma_{j,t}^{n+1}\}_{\forall j,t}$.
- Computation of Likelihood: Obtain $\xi(\theta)$ and $\eta(\theta)$ from recovered mean-utilities δ and ζ, and form likelihood $\mathcal{L}(\theta)$.

[82] Microsoft officially announced the Xbox on March 10, 2000 and Nintendo announced the GC on August 25, 2000, although their existence was rumored for months prior. As often is the case, console manufacturers announce the upcoming release of a new console far in advance to drum up support from software developers and interest from consumers.

[83] This is equivalent to removing the initial values of $\xi_{j,t}$ for the PS2 from the likelihood function specified in the next section.

Details are as follows.

Hardware Adoption

For any inventory state ι, note that the values $\{\delta_{i,j,t}\}_{j\in J_t}$ are sufficient to compute the expected probabilities that a consumer, prior to realizing $\epsilon_{i,t}$, will purchase any hardware console at time t:

$$\hat{s}_{i,t}(\{\delta_{i,j,t}\}_{j\in J_t}, m(t), \beta) = \frac{\exp(\delta_{i,t})}{\exp(\delta_{i,t}) + \exp(\beta E[EV_i(\delta_{i,t+1}, m(t+1))|\delta_{i,t})} \tag{24}$$

as well as the probability a consumer purchases a particular hardware platform j conditional on purchasing any platform:

$$\hat{s}_{i,j|t}(\{\delta_{i,j,t}\}_{j\in J_t}) = \frac{\exp(\delta_{i,j,t})}{\exp(\delta_{i,t})} \tag{25}$$

(These are the standard "logit" closed form expressions derived from integrating over the extreme value errors $\epsilon_{i,t}$). Thus, provided the values of $\{\delta_{i,j,t}\}_{j\in J_t}$ for each consumer i, aggregation over (24) and (25) yields the total predicted share of consumers (who have not yet purchased a console) that purchases console j at time t:

$$\hat{s}_{j,t}(\{\delta_{i,j,t}\}_{j\in J_t}, m(t), \beta, \alpha, \Sigma) = \int_{\alpha^p,\alpha^\gamma} \hat{s}_{it}(\cdot)\hat{s}_{ij|t}(\cdot)dP_t(\alpha^p, \alpha^\gamma) \tag{26}$$

where $dP_t(\cdot)$ represents the distribution of consumer types who have not yet purchased a hardware system at time t, and consumer heterogeneity is parameterized by mean α and variance Σ. The distribution of remaining consumers dP_t changes over time according to the population of consumers who have purchased in previous periods; this is one of the primary ways the demand system generates interdependence over time.

Define the mean utility for the "mean" consumer ($i = 0$) of hardware platform j at time t and inventory state $\iota = 0$ as

$$\delta_{j,t,0} \equiv \alpha^x x_{j,t} - \alpha_0^p p_{j,t} + \Gamma_{j,t,0}(\alpha_0^\gamma, \alpha_0^p) + \xi_{j,t} \tag{27}$$

For a fixed θ_1, $\{\Gamma_{j,t,\iota}\}_{\forall j,t,\iota}$, and $\{\lambda_{r,0}^j\}_{\forall r,j}$, the BLP contraction mapping

$$\delta_{j,t,0}^n = \delta_{j,t,0}^{n-1} + \psi\left(\ln(s_{j,t}^o) - \ln(\hat{s}_{j,t}(\delta_{j,t,0}^{n-1}))\right) \tag{28}$$

is used to obtain to mean console qualities $\delta_{j,t,0}(\cdot)$, where s^o denotes the observed share of potential buyers who purchase console j at time t and $\psi \in (0,1)$ is a "tuning" parameter. For the following discussion, I will omit arguments Γ, β, θ, and $m(t)$ whenever it is possible to do so without confusion.

- At each stage of the mapping, implied market shares $\hat{s}_{j,t}(\delta_{j,t}^{n-1})$ are computed according to:

$$\hat{s}_{j,t}(\delta_{j,t}^{n-1}) = \sum_{\iota=0}^{6}\sum_{i=1}^{R} \lambda_{i,t}\hat{s}_{i,t,\iota}(\delta_{i,t,\iota};\theta)\hat{s}_{i,j,\iota|t} \tag{29}$$

  Recall $\hat{s}_{i,t,\iota}$ is the number of consumers of type i with inventory ι predicted to purchase any console at time t, and $\hat{s}_{i,j,\iota|t}$ are the fraction of those consumers who choose console j.

  - To obtain $\hat{s}_{i,t,\iota}(\delta_{i,t,\iota})$ and $\hat{s}_{i,j,\iota|t}$, the consumer dynamic optimization problem for a given $\delta_{j,t}^{n-1}$ and for every consumer type i and inventory state ι must be solved. Noting $\delta_{i,j,t,\iota} = \delta_{j,t,0} - (\alpha_i^p - \alpha_0^p)p_{j,t} + \Gamma_{j,t}(\alpha_i^\gamma, \alpha_i^p; I_\iota) - (\Gamma_{j,t}(\alpha_0^\gamma, \alpha_0^p; I_0))$, the transition kernel according to the regression in (16) is first updated. I assume there is a finite horizon $\bar{T}$ at which point $\{\delta_{j,\bar{T},0}\}_{\forall j}$ decays to 0, and simulate forward 50 sample paths to compute the expected value function $EV_i(\{\delta_{i,j,t}(I_\iota)\})$ given by (23).[84] In practice, I assume this horizon occurs in January 2006, 3 months after the end

[84] I also explored using a discredited state space with a non-uniform grid (concentrating points in areas that are more likely to be visited), simplical interpolation, and standard value function iteration for convergence. Even when using a relatively sparse grid of 20 points in each direction for the δ_{ijt} terms, the state space is of size $8 \times 12 \times 20^3 = 8 \times 10^5$.

of the data sample; however, results did not change significantly when the horizon was extended by 1 year.

– To compute λ, first shares $\{\lambda_{i,0,0}\}_{\forall i}$ at $t = 0$ are computed from the distribution implied by θ_1, and then each future period is computed by updating the distribution of consumers remaining on the market as follows:

$$\lambda_{i,t+1,\iota} = \frac{(1 - \hat{s}_{i,t,\iota})\lambda_{i,t,s} + \sum_{\iota' \in \mathbf{I}^-(\iota)} \lambda_{i,t,\iota'} \hat{s}_{i,t,\iota'} \hat{s}_{i,\iota \setminus \iota', \iota' | t}}{\sum_{i,\iota'' \neq 7} \lambda_{i,t+1,\iota''}} \tag{30}$$

where $\mathbf{I}^-(\iota)$ is the set of inventory states that can "reach" inventory state ι – e.g., differ only by having one fewer console. In other words, the share of consumers with inventory ι at time $t+1$ are simply those that did not purchase a new console at time t (the first term of the numerator) plus those in state ι' at time t who purchase console j, where j is the only difference between ι and ι'. To account for the growth in total market size (i.e., more television households are present in each period), I assume that new households are distributed across consumer types according to their initial distribution.

Software Adoption

As on the hardware side, note the mean utility for consumer i, $\zeta_{i,k,t}$, is a sufficient statistic in determining whether or not consumer i purchases software title k in a given period. The model implies that the share of people who purchase title k is given by:

$$s_{j,k,t} = \int \frac{\exp(\zeta_{ikt})}{\exp(\zeta_{ikt}) + \exp(\beta E[EW_i(\zeta_{i,k,t+1}|\zeta_{ikt})])} dP_{j,k,t}(\alpha^p, \alpha^\gamma) \tag{31}$$

where $P_{j,k,t}(\cdot)$ is distribution of consumer characteristics for those who own platform j but have not bought software k at time t; it is a function of the hardware adoption decision for all consumers in periods $\tau \leq t$. It is the evolution of this distribution over time and across platforms that necessitates the joint estimation of software and hardware demand.

To obtain a starting value of $\Gamma_{j,t,\iota}$ for the hardware adoption side, I first assume $\{\lambda^j_{i,t}\}_{\forall t} = \lambda^j_{i,0}$ – i.e., the entry distribution of consumer types on each hardware platform is stationary across time. For a given $\theta_1, \{\lambda^j_{i,t}\}_{\forall i,j,t}$, the software side proceeds in a parallel fashion to the hardware adoption side. For each console j, the same BLP contraction mapping is used to recover mean software qualities $\zeta_{j,k,t}$:

$$\zeta^n_{j,k,t}(\theta_1, \{\lambda^j_{i,t}\}) = \zeta^{n-1}_{j,k,t} + \ln(s^o_{j,k,t}) - \ln(\hat{s}_{j,k,t}(\zeta^{n-1}_{j,k,t}))$$

- Implied market shares $\hat{s}_{j,k,t}(\zeta^{n-1}_{j,k,t})$ are computed as in the hardware side (except now there are only two inventory states {0,1}), where the initial base of consumers who have not purchased a title is given by the distribution of consumers on a given console at the time of the title's release, and each future period's potential market size is updated accordingly. Again, the consumer dynamic optimization problem given by (3.2) for a given $\zeta^{n-1}_{j,k,t}$ is solved for every consumer type i, where one can updated $\zeta_{i,j,k,t} = \zeta_{j,k,t} - (\alpha^p_i - \alpha^p_0)p_{k,t} + \alpha^\gamma_i$. I discretize the state space into a uniform grid with 201×12 gridpoints, and employ Halton sequences for random draws on the evolution of $\zeta_{i,j,k,t}$, simple linear interpolation, and standard value function iteration for convergence. At each stage, transition kernel is updated according to the regression in (17).

Initial results were similar to those obtained using a finite horizon, but this method did not scale well when the number of consumer types increased and as a result became too computationally unwieldy. Attempts at using polynomial approximations with shape-preserving splines were also unsatisfactory; similar to the issue raised in Section 9 of Benitez-Silva, Hall, Hitsch, Pauletto, and Rust (2000), polynomial approximation routines do not capture the location of the kink in the value function of a consumer's optimization problem accurately. As this value determines whether or not the consumer purchases, it is of central importance in computing the likelihood function, and thus these inaccuracies render the approximations unsuitable despite their computational advantages.

Once the expected value function is computed for each software title, consumer type, and time period, Γ is updated according to (10).

Recovery of ξ and η

The hardware and software adoption algorithms are repeated until convergence on $\{\lambda_{i,t}^{j}\}_{\forall i,j}$ and $\{\Gamma_{j,t,\iota}\}_{\forall j,t,\iota}$ is obtained. This yields $\{\delta_{j,t,0}(\theta_1)\}_{\forall j,t}$ and $\{\zeta_{j,k,t}(\theta_1)\}_{\forall j,k,t}$.

I next recover $\xi_{j,t}(\theta_1)$ as the residual from the regression of $(\delta_{j,t}(\theta_1) + \alpha_0^p p_{j,t} - \Gamma_{j,t,0}) - \rho^{hw}(\delta_{j,t-1}(\theta_1) + \alpha_0^p p_{j,t-1} - \Gamma_{j,t-1,0})$ on $(x_{j,t} - \rho^{hw} x_{j,t-1})$. Similarly, $\eta_{j,k,t}(\theta_1)$ can be recovered from running the appropriate regression on $\zeta_{k,t}$. The use of MLE and normally distributed errors allows for expressing θ_2 as a linear function of θ_1 in this manner; i.e., θ_2 can be "concentrated" out and a non-linear search is conducted only over θ_1. To estimate the software-title fixed effects, of which there are over a thousand, a partitioned regression is used.

Finally, the log-likelihood provided by (13) can be evaluated from using the residuals defined in (12).

A.4 Fit of Model

Figure 4 plots the predicted values of $\{\delta_{j,t}\}_{\forall j,t}$ for the mean consumer. For this and all following figures, I use the results from the full demand model given by specification (v) in table 3. Consumers' expectations of hardware mean-utilities $\{\delta_{j,t}\}_{\forall j,t}$ are assumed to be a function of the previous values across all consoles as well as the time of year: as evident, values for each console do seem to track each other, and each is drastically affected by seasonality with a large increase occurring during the holiday seasons. To determine whether conditioning only on these past variables, as given by the parameterization in (16), provides a reasonable approximation of consumer expectations, figure 5 plots the error in the realized value of $\delta_{j,t+1}$ from the expected value implied by the estimated transition process $F_i(\{\delta_{j,t}\}_{\forall j}, m(t))$. For the most part, these predicted errors are relatively small, with their variance on average less than 15% of the variance in the actual change in $\delta_{j,t+1} - \delta_{j,t}$. There does not seem to be any persistent correlation or time trends. I also computed these errors without explicitly accounting for seasonality effects, and errors were substantially larger in magnitude – nearly quadrupling in variance – indicating once again the importance of controlling for the time of year.

From the demand estimates, there is significant persistence in hardware and software unobservables with ρ^{hw} estimated to be .65 and ρ^{sw} to be .81. A value of ρ^{hw} at .65 indicates the degree of variance in $\xi_{j,t}$ explained by $\xi_{j,t-1}$ is $.65^2 \approx 42\%$, which indicates there is relatively a large amount of unexplained variation across periods in the hardware unobservable characteristics. However, this variation in ξ given by $\nu_{j,t}^{hw}$ only comprises $10-15\%$ of the total variance in $\delta_{j,t}$ across consoles, and thus does not necessarily indicate lack of explanatory power on the part of the model.

The key assumption used for inference is 3.5, which states the evolution of ν^{hw} and ν^{sw} in unobserved product characteristics ξ and η is independent of each other and the change in observed characteristics at each point in time. Figure 6 plots the implied values of ν^{hw} for each hardware device. These values also appear to be uncorrelated across hardware platforms and across time, something statistical tests cannot reject. Thus, there do not seem to be any significant common shocks across platforms to hardware unobservable characteristics that are not already accounted for by changes in the observed characteristics $x_{j,t}$.

I next examine the predicted number of consumers who multihome and purchase multiple consoles. Although the data indicates 53.2M sixth-gen consoles were sold, the model predicts 47M households actually purchased one or more consoles – i.e., 12% of households are predicted to purchase more than 1 console, in which only a very small number ($< 1\%$) purchase all 3. Nielsen Media estimates that at the end of 2005, 43M console households in the US.[85] Although the model slightly underestimates the amount of multihoming if the Nielsen data is accurate, it does provide substantial correction for the possibility consumers purchase multiple devices. The model also indicates that different consoles exhibit different propensities for multihoming: whereas 15% and 24% percent of PS2 and Xbox owners are predicted to own more than 1 console, over 32% of GC owners do so.

[85] "The State of the Console," *The Nielsen Company*, 2007.

A.5 Robustness Exercises

Alternate Formulation for Future Software Utility

This section details an alternative specification for future software utility, $\Lambda_{j,t}(\cdot)$. I assume that consumers form expectations over future software availability in a two-step procedure. First, every consumer has the same expectations over the *number* of titles released in each month, which is itself is a function only of the number of current software titles available and the time period (for age and month effects); secondly, each consumer perceives the expected mean utility of each new title $\zeta_{i,k,t}, k \in \mathbf{K}_{j,R(t+1)}$ is independently drawn from a distribution that is consistent with the observed ζ's of titles released in a given month.[86]

Assumption A.2. *Let $q_{j,t} \equiv |\boldsymbol{K}_{j,R(t)}|$, the number of new software titles released for platform j at time t, and let $Q_{j,t} \equiv |\boldsymbol{K}_{j,t}|$ be the number of total software titles currently available (i.e., released during and before period t). Consumers perceive the distribution of $q_{j,t}$ is a function only of $Q_{j,t}$, $m(t)$, and the platform j's age $a_{j,t}$:*

$$H(q_{j,t+1}|\Omega_t) = H_j(q_{j,t+1}|Q_{j,t}, m(t+1), a_{j,t+1})$$

Additionally, the probability a future title has quality $\zeta_{i,k,t}$ is a function only of the platform and month in which the title will be released:

$$J(\zeta_{i,k,t+\tau}|\Omega_t) = J_{i,j}(\zeta_{i,k,t+\tau}|m(t+\tau)) \; \forall \; k \in \mathbf{K}_{j,R(t+\tau)}$$

Clearly, the number of titles released in a future period depends on more factors than just the number of titles currently available, most notably the installed base of a given console. Nonetheless, there are two reasons why assumption A.2 may not be problematic: (i) the total number of software titles currently available itself is a proxy for the proclivity of further software development as it should in turn reflect the size of the installed base;[87] (ii) most consumers do not have accurate information as to a console's current installed base, but do know the number of titles currently available (given by simply looking at store shelves).

A consumer thus perceives future software utility of a platform as a series of iterated expectations over both the number ($q_{j,t}$) and quality ($\zeta_{i,k,t}$) of software titles to be released in all future periods:

$$\Lambda_{j,t}(\alpha_i^{\gamma}, \alpha_i^{p}) = E_t\left[\sum_{k=1}^{q_{j,t+1}} \zeta_{i,k,t+1} + E_{t+1}\big[\sum_{k=1}^{q_{j,t+2}} \zeta_{i,k,t} + E_{t+2}[\cdots]\big|Q_{j,t}, q_{j,t+1}\big]\big|Q_{j,t}\right]$$

Computation: Before the main routine is begun, I first estimate "offline" consumers' beliefs over the number of software titles released in future months, which is given by the distribution function $H(\cdot)$ in assumption A.2. A natural candidate would be to use a poisson distribution. However, since there is overdispersion in the data, a negative binomial fixed effects model similar to that proposed in Hausman, Hall, and Griliches (1984) is more suitable. I thus assume $h(\cdot)$ has the density of a negative binomial, but of a specification slightly different than theirs.[88] The mean of the negative binomial is μ and the variance is $\mu + \alpha\mu^2$. I parameterize the conditional mean as $\ln(\mu) = x_{j,t}^{\mu}\beta_\mu$ where $x_{j,t}^{\mu}$ contains both $Q_{j,t}$ as well month dummies and platform age terms, and the dispersion term is parameterized as $\ln(\alpha) = x_j^{\alpha}\beta_\alpha$, where x_j^{α} are platform dummies. For estimation, I use the number of titles released for platforms which were released

[86] This can easily be extended to allow consumers to have adaptive expectations; e.g., expected software utilities are drawn from the observed distribution of only those titles released prior to time t.

[87] Simple tests reject the hypothesis that installed base and software release numbers do not Granger-cause the other.

[88] Cameron and Trivedi (1998) refer to the specification used by Hausman, Hall, and Griliches (1984) as an NB1 model, whereas I use what they refer to as the NB2 specification:

$$h(q|\mu,\alpha) = \frac{\mathbf{\Gamma}(q+\alpha^{-1})}{\mathbf{\Gamma}(q+1)\mathbf{\Gamma}(\alpha^{-1})}\left(\frac{\alpha^{-1}}{\alpha^{-1}+\mu}\right)^{\alpha^{-1}}\left(\frac{\mu}{\alpha^{-1}+\mu}\right)^{q}$$

where $\mathbf{\Gamma}$ is the gamma function (not to be confused with the definition of software utility $\Gamma(\cdot)_{j,t}$ used in this paper): $\mathbf{\Gamma}(a) = \int_0^\infty e^{-t}t^{a-1}dt$, $a > 0$

in the sixth as well as fifth-generation, and extend out the sample to April, 2007.[89] Once this model is estimated, there is still the issue of forming consumer expectations at time t over the path of software titles in all future periods. This integral is intractable, but can be approximated via forward simulation. For each console and every t, I first take a draw on $q_{j,t+1}$ conditional on observed $Q_{j,t}$, and then proceed forward by updating $\tilde{Q}_{j,t+1} = Q_{j,t} + \tilde{q}_{j,t+1}$, where $\tilde{q}_{j,t+1}$ is the value previously drawn. Updating to a terminal value and creating multiple sample paths allows for the approximation of $\{E_t[\{q_{j,t+\tau}\}_{\forall\tau>0,\forall j}]\}_{\forall t}$.

Table 14 presents the results of the negative binomial regression on the number of software titles released in the next month as a function of the total number of titles currently available and age of the console. The more titles currently available indicates the higher likelihood of titles being released in the future, whereas age negatively affects this probability. Seasonality plays a significant role, with most titles likely to be released in the months prior to the holiday season. Finally, there is overdispersion predicted in the data, and α is predicted to be the largest for the PS2, which had the most titles released in the sample, and the smallest for the GC, which had the least number of titles.

One question is the fit of this model in predicting the future number of titles released. Figure 7 graphs the number of titles released for each console the 5% and 95% percentile of the predicted negative binomial distribution. With only a few exceptions, the number of titles released for all three consoles fall within the interval. The larger variance of titles released during the peak months and the differing variances across consoles is captured by the model. Though only the dispersion fixed-effects for the sixth generation consoles are reported, including the data from the previous generation of consoles helped match the observed frequency of software releases to the confidence interval predicted by the model.

B Software Network Formation Game: Further Details

B.1 Assumptions and Computation of Profits and Costs

I make several assumptions governing how software firms form their expectations over the profitability of developing for different sets of consoles, which includes how other software titles might respond to its choice.

Assumption B.1. $E_{R(k)-\tau}[p_{k,R(k)}] = p_{k,R(k)}$ *and* $E_{t-\tau}[\zeta_{k,R(k)}] = \zeta_{k,R(k)}$.

This assumption states that every software firm knows the initial price and quality of its title at the time of release. As almost all titles are released at starting price points of \$49.99, \$39.99, and \$29.99, a firm typically knows the pricing entry point of its title before it goes to market. During development, the firm should also have knowledge as to the title's eventual quality, as proxied by budget and early feedback from beta tests.

Assumption B.2. *Software firms perceive their price will follow a first-order Markov process governed by* $p_{k,t'} = F_p(p_{k,t'-1}, m(t'))\ \forall\ t' > R(k)$

Table 10 provides an OLS regression of software prices on lagged prices, and shows that nearly all of the variation in prices is contained in the previous period's value. Additionally, although the number of titles sold in the previous period is also a significant factor in determining next period's prices, its inclusion in this and the subsequent exercises did not significantly affect results; this is unsurprising given the limited improvement in the R^2 of the initial regression.

Assumption B.3. *Firms perceive their mean-quality for all consumer types* $\zeta_{i,k,t}$ *can be summarized by the first-order Markov process given by* (9).

Assumption B.4. *Firms perceive hardware mean-quality for all consumer types* $\delta_{i,j,t}$ *can be summarized by the first-order Markov process given by* (22).

Assumptions B.3 and B.4 deal with the expectations a software firm has over the evolution of its own quality and the quality of each hardware platform, and imply that firms share the same expectations as

[89] The number of software titles released for each of the three sixth-generation consoles from November 2005 to April 2007 was collected from `http://www.gamespot.com`, an online videogame website.

consumers do over the evolution of products.[90] When the environment does not significantly change (which I will assume to be the case when only unilateral deviations are considered when computing porting costs), these transition processes are the same as those already estimated from the demand system. Otherwise, as will be the case when recomputing the contracting decisions for all software titles in the counterfactual regimes, these transition probabilities will be re-estimated. Software title k is thus assumed to take the transition probabilities over $\{\delta_{i,j,t}\}_{\forall i,j,t}$ as given, and sees its own action as only affecting the state at time t but not other agents' perceptions of the transition probabilities. Since title k is only affected by the strategic actions of other titles only in so far as they affect the potential number of consumers on each console – i.e., through $\{\delta_{i,j,t}\}_{\forall i,j,t}$ – assumption B.4 allows for a title k to internalize how its choice of strategy affects the actions of future software titles without explicitly needing to account for those decisions.

Finally, when predicting the values of $\zeta_{i,k,t}$ for title on a platform that it was not observed to have joined, parameter estimates from the demand side, α^y, are used. The title fixed effect is adjusted according to platform fixed effects, separately estimated from a regression of software fixed effects on platform dummies for those titles that multihomed.

I rewrite (18) as:

$$E[\pi_k(\mathbf{s}_k;\theta_C)|\Omega_{r_k-\tau}] = \left(\sum_{t=R(k)}^{T} \beta^{\tau+t-R(k)} \sum_{j\in\mathbf{s}_k} E[M_{j,k,t}s_{j,k,t}((1-rmkup_t)p_{j,k,t}-mc_j)]\right) - C_k(\mathbf{s}_k;\theta_C)$$

where now $Q_{j,k,t}$ has been broken into two different components: $s_{j,k,t}$, which represents the share of consumers who purchase title k, and $M_{j,k,t}$, which represents the number of consumers on platform j who have not yet purchased title k. $s_{j,k,t}$ is defined in (31), and is solely a function of $\zeta_{k,t}$ and the distribution of consumer types onboard platform j who have not yet purchased the title. If $IB_{j,t}$ is the number of consumers who own console j at time t, and $IB_{j,k,t}$ the number of consumers who own title k on platform j, then $M_{j,k,t} = IB_{j,t} - IB^k_{j,t}$, where $IB^k_{j,t} = IB^k_{j,t-1} + M_{j,k,t-1}s_{j,k,t-1}$. From the demand side, recall a sufficient statistic for determining $IB_{j,t}$ is $IB_{j,t-1}$ and $\{\delta_{j,t}\}_{\forall j}$.[91]

To form the first part of $E[\pi_k(\mathbf{s}_k);\theta_C]$, only expected values of $\{\{\delta_{j,t}\},\{\zeta_{j,k,t}\},\{p_{j,k,t}\}\}_{\forall j,\forall t>R(k)-\tau}$ are first required. I obtain these using a simulated frequency approach as in Pakes (1986): multiple sample paths of these variables are created via forward simulation using the estimated transition processes from the demand system (given by (16), (17), and table 10), and the appropriate quantities $M_{j,k,t}$ and $s_{j,k,t}$ are calculated at each point in time, again from the demand system. At release date r_k, the predicted hardware mean utilities $\{\delta_{j,t}\}$ are increased by the amount software k contributes to each platform it joins, as determined by its choice of strategy $\mathbf{s}_k$.

C Proofs

C.1 Proof of Proposition 3.6

The proof follows closely Chanda (1954) and Bradly and Gart (1962), and provides necessary modifications where needed.[92] The assumptions required for the proof are as follows:

Assumption C.1. *Let $\theta = (\theta_1,\ldots,\theta_k)$ and $i \in \{hw, sw\}$.*

[90] Note that ζ includes price, and yet is assumed to evolve in an independent process from price itself. To address this concern, I also estimate an alternative specification whereby I assume software mean quality net of price (i.e., $\widetilde{\zeta_{i,k,t}} = \zeta_{i,k,t} + \alpha_i^{p,sw}p_{k,t}\ \forall\ i$) evolves according to a Markov process, and proceed in a similar fashion (whereby $\zeta_{k,t}$ is constructed in each period from the separate evolution in $\widetilde{\zeta_{k,t}}$ and $p_{k,t}$). Results do not change in any significant way.

[91] Although for notational simplicity I have omitted discussing the distribution of consumer heterogeneity and inventory states, these concerns are not ignored in estimation.

[92] NB: there is an error in Chanda (1954)'s proof of uniqueness which is addressed and fixed in Tarone and Gruenhage (1975) (which also has a corrigenda in 1979).

i. For almost all $x \in \Re$ and for all $\theta \in \Theta$

$$\frac{\partial \ln f^i}{\partial \theta_r}, \quad \frac{\partial^2 \ln f^i}{\partial \theta_r \theta_s} \quad \textit{and} \quad \frac{\partial^3 \ln f^i}{\partial \theta_r \theta_s \theta_t} \qquad r, s, t = 1, \ldots, k$$

exist.

ii. For almost all $x \in \Re$ and for all $\theta \in \Theta$,

$$\left| \frac{\partial f^i}{\partial \theta_r} \right| < F_{ir}(x), \quad \left| \frac{\partial^2 \ln f^i}{\partial \theta_r \theta_s} \right| < F_{irs}(x) \quad \textit{and} \quad \left| \frac{\partial^3 f^i}{\partial \theta_r \theta_s \theta_t} \right| < H_{irst} \qquad r, s, t = 1, \ldots, k$$

where $F_{ir}(x)$ and $F_{irs}(x)$ are integrable over $\Re$, and $E[H_{irst}(x)] < M^i$, where M^i are finite positive constants.

iii. For all $\theta \in \Theta$, the matrix $\mathbf{J} = [J_{rs}(\theta)]$ with

$$J_{rs}(\theta) = \left[\sum_{j=1}^{J} \mu_j^{hw} \int \left(\frac{\partial \ln f^{hw}}{\partial \theta_r} \frac{\partial \ln f^{hw}}{\partial \theta_s} \right) f^{hw} dx + \mu_j^{sw} \int \left(\frac{\partial \ln f^{sw}}{\partial \theta_r} \frac{\partial \ln f^{sw}}{\partial \theta_s} \right) f^{sw} dx \right]$$

is positive definite with finite determinant.

These are different from, and in a sense stronger than, those used in Wald (1949) for his general proof on the consistency of MLE. The assumptions here, however, are easier to check and confirm.

Recall the log-likelihood function to be maximized is

$$\mathcal{L}(\theta) = \sum_{j=1}^{J} \sum_{t=r_j+1}^{T} \left(f^{hw}(\nu_{j,t}^{hw}(\theta); \theta) + \sum_{k=1}^{K_{j,t}} f^{sw}(\nu_{j,k,t}^{sw}(\theta); \theta) \right)$$

and $n_j^{hw} = \sum_{t=r_j}^{T} 1$ and $n_j^{sw} = \sum_{t=r_k}^{T} \sum_{k=1}^{K_{j,t}} 1$ be the number of observations for platform j, and $N = \sum_{j=1}^{J} (n_j^{hw} + n_j^{sw})$ be the total number of error observations. Let θ^0 represent the true value of θ.

Consistency: Let $\theta^0 \in \Theta$ be the true value of the parameter vector θ to be estimated. Take the following Taylor expansion for $i \in \{hw, sw\}$:

$$\frac{\partial \ln f^i}{\partial \theta_r} = \left(\frac{\partial \ln f^i}{\partial \theta_r} \right)_{\theta=\theta_0} + \sum_{s=1}^{k} (\theta_s - \theta_s^0) \left(\frac{\partial^2 \ln f^i}{\partial \theta_r \partial \theta_s} \right)_{\theta=\theta_0} + \frac{1}{2} \sum_{s,q=1}^{k} (\theta_s - \theta_s^0)(\theta_q - \theta_q^0) \left(\frac{\partial^3 \ln f^i}{\partial \theta_r \partial \theta_s \partial \theta_q} \right)_{\theta=\theta'}$$

where $\theta' = \theta'(x)$ is a value that depends on x but for all x lies within the hyper-cell containing $\theta - \theta^0$ as its diagonal. Summing these Taylor expansions across i, j, t, k and dividing by N yields:

$$\frac{1}{N} \frac{\partial \mathcal{L}}{\partial \theta_r} \equiv L_r(\theta) = L_r(\theta^0) - \sum_{s=1}^{k} \delta_s L_{rs}(\theta^0) + \frac{1}{2} \sum_{s,t=1}^{k} \delta_s \delta_t L_{rst} \tag{32}$$

where

$$
\begin{aligned}
\delta_s &= \theta_s - \theta_s^0 \\
L_r(\theta) &= \sum_{j=1}^{J}\left(\mu_j^{hw}\sum_{t=r_j}^{T}\frac{1}{n_j^{hw}}\frac{\partial \ln f_{j,t}^{hw}}{\partial\theta_r} + \mu_j^{sw}\sum_{t=r_k}^{T}\sum_{k=1}^{K_{j,t}}\frac{1}{n_j^{sw}}\frac{\partial \ln f_{j,t}^{sw}}{\partial\theta_r}\right) \\
L_{rs}(\theta) &= \sum_{j=1}^{J}\left(\mu_j^{hw}\sum_{t=r_j}^{T}(-1)\frac{1}{n_j^{hw}}\frac{\partial^2 \ln f_{j,t}^{hw}}{\partial\theta_r\partial\theta_s} + \mu_j^{sw}\sum_{t=r_k}^{T}\sum_{k=1}^{K_{j,t}}(-1)\frac{1}{n_j^{sw}}\frac{\partial^2 \ln f_{j,t}^{sw}}{\partial\theta_r\partial\theta_s}\right) \\
L_{rsq} &= \sum_{j=1}^{J}\left(\mu_j^{hw}\sum_{t=r_j}^{T}\frac{1}{n_j^{hw}}\left.\frac{\partial^3 \ln f_{j,t}^{hw}}{\partial\theta_r\partial\theta_s\partial\theta_q}\right|_{\theta=\theta'} + \mu_j^{sw}\sum_{t=r_k}^{T}\sum_{k=1}^{K_{j,t}}\frac{1}{n_j^{sw}}\left.\frac{\partial^3 \ln f_{j,t}^{sw}}{\partial\theta_r\partial\theta_s\partial\theta_q}\right|_{\theta=\theta'}\right)
\end{aligned}
$$

From C.1 it follows that

$$
\begin{aligned}
L_r(\theta^0) &\to_p 0 \\
L_{rs}(\theta^0) &\to_p J_{rs}(\theta^0) \\
|L_{rsq}| &< \sum_{j=1}^{J}\mu_j^{hw}M^{hw} + \mu_j^{sw}M^{sw}
\end{aligned}
$$

With both f^{hw} and f^{sw} to consider, both Khintchine's Theorem (Rao (1973), 2.c.ii) and Slutsky's Theorem are required for convergence results. E.g., since $E[\partial \ln f_{j,t}^i/\partial\theta_r] = 0$, Khintchine's Theorem proves $\sum_{t=r_j}^{T}(1/n_j^i)(\partial \ln f_{j,t}^i/\theta_r) \to_p 0$; Slutsky's Theorem handles the sum of two convergent sequences and leads to the result $L_r(\theta) \to_p 0$. The same arguments apply for the last two results.

The remainder of the consistency proof follows Chanda (1954) using the same notation, and is omitted here.

Asymptotic Normality: Following Bradly and Gart (1962), substitute $\hat\theta$ into (32), note $L_r(\hat{(\theta)}) = 0$, and rearrange to obtain:

$$
L_r(\theta^0) = \sum_s(\hat\theta_s - \theta_s^0)L_{rs}(\theta^0) - \frac{1}{2}\sum_{s,q=1}^{k}(\hat\theta_s - \theta_s^0)(\hat\theta_q - \theta_q^0)L_{rsq}
$$

for $r = 1, \ldots, k$, or in matrix notation:

$$
(\mathbf{L}+\mathbf{G})(\hat\theta' - \theta^{0'}) = \left[\sum_{j=1}^{J}\left(\mu_j^{hw}\sum_{t=r_j}^{T}\frac{1}{n_j^{hw}}\left.\frac{\partial \ln f_{j,t}^{hw}}{\partial\theta_r}\right|_{\theta=\theta^0} + \mu_j^{sw}\sum_{t=r_k}^{T}\sum_{k=1}^{K_{j,t}}\frac{1}{n_j^{sw}}\left.\frac{\partial \ln f_{j,t}^{sw}}{\partial\theta_r}\right|_{\theta=\theta^0}\right)\right]' \tag{33}
$$

where $\mathbf{L} = [L_{rs}(\theta^0)]$ and $\mathbf{G} = [(-1/2)\sum_{q=1}^{k}(\hat\theta_q - \theta_q^0)L_{rsq}]$ are k-square symmetric matrices. In the proof of consistency, I have shown $\mathbf{L} \to_p \mathbf{J}_0$ (which is positive definite, by assumption), and also can show $G \to_p 0$ since $\hat\theta \to_p \theta^0$ and $|L_{rsq}|$ is bounded. Thus, for large T, $\mathbf{L}+\mathbf{G}$ may be inverted and by Slutsky's Theorem, $\mathbf{L}+\mathbf{G} \to_p \mathbf{J}_0$ and $(\mathbf{L}+\mathbf{G})^{-1} \to_p \mathbf{J}_0^{-1}$. Rewriting (33) yields:

$$
\sqrt{N}(\hat\theta' - \theta^{0'}) = (\mathbf{L}+\mathbf{G})^{-1}\left[\sum_{j=1}^{J}\left(\mu_j^{hw}\sum_{t=r_j}^{T}\frac{1}{n_j^{hw}}\left.\frac{\partial \ln f_{j,t}^{hw}}{\partial\theta_r}\right|_{\theta=\theta^0} + \mu_j^{sw}\sum_{t=r_k}^{T}\sum_{k=1}^{K_{j,t}}\frac{1}{n_j^{sw}}\left.\frac{\partial \ln f_{j,t}^{sw}}{\partial\theta_r}\right|_{\theta=\theta^0}\right)\right]'
$$

where the last term on the right, by the multivariate central limit theorem (Rao (1973), 2.c.iv), tends to a multivariate normal distribution with mean zero and variance covariance $\mathbf{J_0}$. It thus follows that the asymptotic distribution of $\sqrt{N}(\hat\theta' - \theta^{0'})$ is multivariate normal with mean zero and variance covariance matrix $\mathbf{J_0}^{-1}\mathbf{J_0}\mathbf{J_0}^{-1} = \mathbf{J_0}^{-1}$ as $T \to \infty$ and μ_j^i constant $\forall i \in \{hw, sw\}, j \in J$.

Table 1: Industry Summary Statistics

	PS2	XBOX	GC	ALL
HARDWARE				
Release Date	Oct. 2000	Nov 2001	Nov 2001	
Months Active	61	48	48	
Price				
Average	$213.46	$185.47	$127.79	
Std. Dev	62.72	47.36	36.50	
Max	316.29	299.97	199.85	
Min	147.37	146.92	92.37	
Quantity (M)				
Average	0.49	0.28	0.20	0.87
Std. Dev	0.44	0.25	0.23	0.86
Max	2.69	1.08	1.16	4.31
Min	0.18	0.08	0.04	0.18
Installed Base (Oct 2005)	30.07	13.32	9.83	53.22
SOFTWARE				
Total Quantity Sold (M)				
Average	0.26	0.16	0.16	0.21
Std. Deviation	0.10	0.07	0.07	0.08
Median	0.50	0.30	0.28	0.41
Max	6.69	4.82	2.96	6.69
Min	0.00	0.00	0.00	0.00
Starting Price				
Avg.	$40.83	$39.63	$39.72	$40.24
Std. Deviation	11.22	10.69	9.83	10.80
Total # Titles Released	1161	749	487	1581
% Exclusive	52.4	33.4	27.5	62.7
% Exclusive, First Party	8.9	8.4	9.2	13.3
% Also on PS2		63.4	67.6	73.4
% Also on XB	40.9		56.3	47.4
% Also on GC	28.3	36.6		30.8
% On all 3 Consoles	21.6	33.5	51.5	15.9
# Titles Released Per Month				
Average	19	13	8	37
Max	54	45	38	127
Min	2	0	0	5

Notes: Summary statistics for the PS2 are for the 61-month period between October 2000 and October 2005; for the other two consoles are for a 48-month period beginning on November 2001.

Table 2: Top 10 Videogame Titles for Each Platform by Quantity Sold

Console	Title	Publisher	Release Date	Exclusive?	Quantity ('000s)
PS2	GTA: Vice City	Take 2	Oct 2002	Yes[a]	6,687
	GTA: San Andreas	Take 2	Oct 2004	Yes[b]	5,797
	GTA 3	Take 2	Oct 2001	Yes[c]	5,588
	Gran Turismo 3: A-Spec	Sony	Jul 2001	Yes	3,764
	Madden NFL 2004	EA	Aug 2003	No	3,446
	Madden NFL 2005	EA	Aug 2004	No	3,199
	Madden NFL 2003	EA	Aug 2002	No	2,728
	Need for Speed: UG	EA	Nov 2003	No	2,533
	Kingdom Hearts	Square	Sep 2002	Yes	2,460
	Medal of Honor: Frontline	EA	May 2002	No	2,325
Xbox	Halo 2	Microsoft	Nov 2004	Yes	4,822
	Halo	Microsoft	Nov 2001	Yes	4,157
	T. Clancy's Splinter Cell	Ubisoft	Nov 2002	No	1,521
	GTA Pack	Take 2	Oct 2003	No	1,379
	Madden NFL 2005	EA	Aug 2004	No	1,223
	Project Gotham Racing	Microsoft	Nov 2001	Yes	1,216
	Star Wars: KOTR	Lucasarts	Jul 2003	Yes	1,210
	ESPN NFL 2K5	Take 2	Jul 2004	No	1,200
	Fable	Microsoft	Sep 2004	Yes	1,125
	T. Clancy's Ghost Recon	Ubisoft	Nov 2002	No	998
GC	Super Smash Bros	Nintendo	Dec 2001	Yes	2,961
	Super Mario Sunshine	Nintendo	Aug 2002	Yes	2,157
	Zelda: Wind Waker	Nintendo	Mar 2003	Yes	2,000
	Mario Kart: Double	Nintendo	Nov 2003	Yes	1,922
	Luigi's Mansion	Nintendo	Nov 2001	Yes	1,748
	Metroid Prime	Nintendo	Nov 2002	Yes	1,312
	Sonic Adventures 2	Sega	Feb 2002	Yes	1,188
	Animal Crossing	Nintendo	Sep 2002	Yes	1,099
	Pokemon Colliseum	Nintendo	Mar 2004	Yes	979
	Mario Party 4	Nintendo	Oct 2002	Yes	979

Notes: Summary statistics for the PS2 are for the 61-month period between October 2000 and October 2005; for the other two consoles are for a 48-month period beginning on November 2001. [a,b,c] PS2 had a window of exclusivity for these three titles: GTA: Vice City and GTA: 3 were not released on the Xbox until 2003 (well after they had both become blockbusters for the PS2), whereas GTA: San Andreas, though initially developed for both Xbox and PS2, was not released for the Xbox until 6 months after the PS2 game's release.

Table 3: Estimated Parameters of Demand System

		No Consumer Heterogeneity						Full Consumer Heterogeneity			
		Static Model Singlehoming (i)		Dynamic Model SingleHoming (ii)		Dynamic Model Multihoming (iii)		Dynamic Model Singlehoming (iv)		Dynamic Model Multihoming (v)	
	Variable	Parameter Estimate	Standard Error	Parameter Estimate	Standard Error	Parameter Estimate	Standard Error	Parameter Estimate	Standard Error	Parameter Estimate	Standard Error
Nonlinear	ρ^{hw}			0.618	0.058	0.649	0.050	0.657	0.047	0.652	0.041
Parameters	ρ^{sw}			0.807	0.001	0.807	0.001	0.807	0.001	0.807	0.001
θ_1	$\alpha_0^{p,sw}$	0.876	0.009	1.310	0.005	1.311	0.005	1.335	0.027	1.378	0.030
	σ_{α_p}							0.075	0.082	0.209	0.072
	σ_{α^γ}							0.983	0.056	0.781	0.092
	σ_ϵ	1.598	0.192	1.435	0.170	1.351	0.176	1.590	0.122	1.748	0.126
	β_γ			0.937	0.034	0.931	0.045	0.908	0.012	0.885	0.026
	D					1.187	0.212			0.110	0.399
	Implied α_0^p	1.400		1.978		1.887		2.274		2.511	
Hardware	d_{PS2}	-14.194	1.899	5.309	0.653	4.880	0.614	8.548	0.577	9.370	0.543
Parameters	d_{XBOX}	-22.857	2.857	0.545	0.603	-0.040	0.571	1.462	0.538	2.889	0.505
α^x	d_{GC}	-19.785	2.896	0.586	0.595	-0.154	0.565	1.383	0.532	2.344	0.499
	age	-0.127	0.021	0.032	0.014	0.015	0.014	-0.030	0.014	-0.044	0.013
	age^2 (10^-3)	0.001	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.001	0.000
	IB$_{\text{PS2}}$	1.127	0.138	-0.166	0.037	-0.161	0.033	-0.159	0.031	-0.165	0.029
	IB$_{\text{XBOX}}$	1.673	0.204	0.091	0.034	0.088	0.031	0.236	0.029	0.166	0.027
	IB$_{\text{GC}}$	1.438	0.210	0.045	0.034	0.047	0.031	0.179	0.029	0.128	0.027
Software	age	-0.254	0.001	-0.197	0.003	-0.198	0.003	-0.188	0.003	-0.190	0.003
Parameters	age^2 (10^-3)	0.002	0.000	0.001	0.000	0.001	0.000	0.001	0.000	0.001	0.000
α^y	IB$_{\text{PS2}}$	0.100	0.010	0.056	0.002	0.056	0.002	0.057	0.002	0.056	0.002
	IB$_{\text{XBOX}}$	0.149	0.013	0.025	0.002	0.025	0.002	0.025	0.002	0.025	0.002
	IB$_{\text{GC}}$	0.309	0.012	0.045	0.003	0.045	0.003	0.045	0.003	0.045	0.003
Log Likelihood		-63486		-27654		-27649		-27617		-27613	
# HW Obs. (n^{hw})		154		154		154		154		154	
# SW Obs. (n^{sw})		57901		57901		57901		57901		57901	

Table 4: Estimated Parameters of Demand System (continued): Month Effects

	Variable	Parameter Estimate	Standard Error		Variable	Parameter Estimate	Standard Error
Hardware	d_{Feb}	0.191	0.065	Software	d_{Feb}	0.090	0.012
All Systems	d_{Mar}	0.078	0.084	Xbox	d_{Mar}	0.101	0.015
	d_{Apr}	-0.266	0.095		d_{Apr}	-0.262	0.018
	d_{May}	-0.533	0.100		d_{May}	-0.334	0.019
	d_{Jun}	-0.085	0.104		d_{Jun}	0.087	0.020
	d_{July}	-0.413	0.105		d_{July}	-0.011	0.020
	d_{Aug}	-0.513	0.104		d_{Aug}	-0.017	0.020
	d_{Sep}	-0.336	0.102		d_{Sep}	0.028	0.019
	d_{Oct}	-0.554	0.097		d_{Oct}	-0.167	0.018
	d_{Nov}	0.031	0.087		d_{Nov}	-0.045	0.016
	d_{Dec}	1.037	0.066		d_{Dec}	1.011	0.012
Software	d_{Feb}	0.086	0.008	Software	d_{Feb}	0.056	0.014
Playstation 2	d_{Mar}	0.102	0.011	Gamecube	d_{Mar}	0.058	0.018
	d_{Apr}	-0.185	0.013		d_{Apr}	-0.275	0.020
	d_{May}	-0.216	0.014		d_{May}	-0.345	0.022
	d_{Jun}	0.162	0.014		d_{Jun}	0.088	0.023
	d_{July}	0.087	0.014		d_{July}	0.018	0.023
	d_{Aug}	0.040	0.014		d_{Aug}	0.049	0.023
	d_{Sep}	0.116	0.014		d_{Sep}	0.048	0.022
	d_{Oct}	-0.095	0.013		d_{Oct}	-0.077	0.021
	d_{Nov}	0.110	0.011		d_{Nov}	0.252	0.018
	d_{Dec}	1.120	0.008		d_{Dec}	1.207	0.014

Notes: Coefficients on month dummies are from model specification (v) in Table 3.

Table 5: Regression of Software Title Fixed Effects

	All Titles		PS2 Only		XBOX Only		GC Only	
Variable	Parameter Estimate	Standard Error	Parameter Estimate	Standard Error	Parameter Estimate	Standard Error	Parameter Estimate	Standard Error
Exclusive, VI	0.626	0.132	0.497	0.183	-0.002	0.261	1.882	0.276
Exclusive, 3rd Party	-0.334	0.078	-0.356	0.100	-0.391	0.165	-0.213	0.196
PS2	-4.574	0.101	-4.536	0.126				
Xbox	-3.962	0.107			-4.112	0.161		
GC	-3.794	0.121					-3.645	0.179
Feb	-0.300	0.165	-0.177	0.256	-0.092	0.305	-0.787	0.323
Mar	-0.113	0.173	-0.177	0.215	0.060	0.371	0.105	0.519
Apr	-0.293	0.217	-0.379	0.293	0.050	0.422	-0.694	0.507
May	-0.068	0.200	0.151	0.304	-0.199	0.352	-0.182	0.419
June	-0.497	0.193	-0.449	0.245	-0.316	0.416	-0.932	0.472
July	-0.225	0.206	-0.740	0.313	0.235	0.367	-0.264	0.426
Aug	0.315	0.151	-0.035	0.261	0.696	0.271	0.049	0.285
Sept	0.265	0.134	0.241	0.189	0.490	0.262	0.054	0.295
Oct	0.532	0.127	0.558	0.184	0.673	0.247	0.319	0.277
Nov	-0.032	0.147	0.021	0.187	0.395	0.353	-0.837	0.345
Dec	0.137	0.220	0.216	0.284	0.244	0.453	-0.270	0.534

Notes: OLS Regression of recovered software fixed effects α_k^y for each software title on dummy variables indicating whether or not it was exclusive, the platform it was released on, and the month of release. Exclusive, VI indicates title was published by a platform provider; Exclusive, other indicates another party published the title.

Table 6: Estimated Hardware Own and Cross-Price Semi-Elasticities

		PS2	XBOX	GC	Outside
Model (i)	PS2	0.157	-0.002	-0.002	-0.279
Static		*(.132, .194)*	*(-.002, -.002)*	*(-.003, -.002)*	*(-.344, -.246)*
No Multihoming	XBOX	-0.001	0.157	-0.001	-0.151
No Heterogeneity		*(-.001, -.001)*	*(.133, .194)*	*(.002, -.001)*	*(-.187, -.128)*
	GC	-0.001	-0.001	0.158	-0.108
		(-.001, -.001)	*(-.001, -.001)*	*(.133, .194)*	*(-.133, -.091)*
Model (iii)	PS2	0.166	-0.033	-0.029	-0.067
Dynamic		*(.131, .205)*	*(-.043, -.024)*	*(-.038, -.021)*	*(-.083, -.052)*
Multihoming	XBOX	-0.009	0.187	-0.011	-0.032
No Heterogeneity		*(-.012, -.007)*	*(.147, .232)*	*(-.014, -.008)*	*(-.040, -.024)*
	GC	-0.007	-0.009	0.192	-0.023
		(-.009, -.005)	*(-.012, -.007)*	*(.151, .239)*	*(-.028, -.017)*
Model (v)	PS2	0.199	-0.049	-0.038	-0.066
Dynamic		*(.178, .199)*	*(-.059, -.037)*	*(-.049, -.025)*	*(-.073, -.056)*
Multihoming	XBOX	-0.013	0.234	-0.016	-0.029
Heterogeneity		*(-.016, -.010)*	*(.209, .250)*	*(-.021, -.011)*	*(-.033, -.024)*
	GC	-0.008	-0.012	0.245	-0.020
		(-.011, -.004)	*(-.016, -.008)*	*(.219, .262)*	*(-.023, -.016)*

Notes: Cell entries i, j, where i indexes row and j indexes column, provides the percentage change in market share with a 10% decrease in the price of console i. 95% confidence intervals are provided in parenthesis below estimates.

Table 7: Estimated Software Own-Price Semi-Elasticities

	DYN	MH	HET	PS2	XBOX	GC
Model (i)	No	No	No	0.181	0.752	0.511
				(.181, .181)	*(.752, .752)*	*(.511, .511)*
Model (ii)	Yes	No	No	0.120	0.409	0.278
				(.120, .120)	*(.409, .409)*	*(.278, .278)*
Model (iii)	Yes	Yes	No	0.097	0.147	0.086
				(.097, .097)	*(.147, .147)*	*(.086, .086)*
Model (iv)	Yes	No	Yes	0.106	0.414	0.263
				(.063, .155)	*(.322, .524)*	*(.218, .318)*
Model (v)	Yes	Yes	Yes	0.110	0.165	0.089
				(.041, .129)	*(.068, .219)*	*(.037, .115)*

Notes: Percentage change in total quantity sold of a top selling title on each console conditional on 10% decrease in the price of that title. The software titles are *Grand Theft Auto III* for the PS2, *Halo* for the Xbox, and *Super Smash Bros.* for the GC. 95% confidence intervals are provided in parenthesis below estimates.

Table 8: Hardware Elasticities from Losing A Top Title

		PS2	XBOX	GC	Outside
Model (i)	PS2	-0.018	0.000	0.000	0.033
Static		*(-.021, -.016)*	*(.000, .000)*	*(.000, .000)*	*(.028, .039)*
No Multihoming	XBOX	0.001	-0.056	0.001	0.055
No Heterogeneity		*(.000, .000)*	*(-.066, -.050)*	*(.000, .000)*	*(.048, .064)*
	GC	0.000	0.000	-0.043	0.030
		(.000, .000)	*(.000, .000)*	*(-.038, -.043)*	*(.035, .030)*
Model (iii)	PS2	-0.024	0.007	0.006	0.009
Dynamic		*(-.029, -.020)*	*(.005, .009)*	*(.005, .08)*	*(.008 .011)*
Multihoming	XBOX	0.008	-0.090	0.009	0.014
No Heterogeneity		*(.006, .009)*	*(-.104 -.076)*	*(.007, .010)*	*(.011, .016)*
	GC	0.003	0.004	-0.064	0.007
		(.003, .004)	*(.003, .005)*	*(-.075, -.053)*	*(.006, .008)*
Model (v)	PS2	-0.022	0.013	0.013	0.005
Dynamic		*(-.029, -.018)*	*(.010, .018)*	*(.008, .014)*	*(.004, .007)*
Multihoming	XBOX	0.014	-0.092	0.022	0.006
Full Heterogeneity		*(.010, .017)*	*(-.116, -.072)*	*(.014, .028)*	*(.004, .008)*
	GC	0.004	0.006	-0.067	0.004
		(.003, .005)	*(.004, .007)*	*(-.085, -.054)*	*(.003, .005)*

Notes: Cell entries i, j, where i indexes row and j indexes column, provides the percentage change in market share of console j upon console i losing a top software title. The software titles are *Grand Theft Auto III* for the PS2, *Halo* for the Xbox, and *Super Smash Bros.* for the GC. 95% confidence intervals are provided in parenthesis below estimates.

Table 9: Hardware Elasticities from Forced Compatibility of Software Titles

	DYN	MH	HET	PS2	XBOX	GC	Outside
Model (i)	No	No	No	2.294	1.039	1.292	-
				(.705, 1.13)	*(-.141, .050)*	*(.150, .224)*	-
Model (ii)	Yes	No	No	1.081	-0.069	0.737	-0.932
				(.835, 1.691)	*(-.520, .257)*	*(-.125, .916)*	*(-.966, -.854)*
Model (iii)	Yes	Yes	No	2.068	-0.668	0.413	-0.962
				(1.823, 2.561)	*(-.843, -.503)*	*(-.834, .894)*	*(-.972, -.940)*
Model (iv)	Yes	No	Yes	0.461	-0.063	0.044	-0.311
				(.345, .581)	*(-.103, -.018)*	*(.000, .106)*	*(-.380, -.257)*
Model (v)	Yes	Yes	Yes	0.745	-0.124	0.168	-0.389
				(.574, 1.10)	*(-.195, -.082)*	*(.018, .270)*	*(-.543, -.329)*

Notes: Percentage change in market share of each console subject to every software title "mulithoming" and joining all three consoles. 95% confidence intervals are provided in parenthesis below estimates.

Table 10: Software Price Regressions

Explanatory Variable	Parameter Estimate	Standard Error	Parameter Estimate	Standard Error	Parameter Estimate	Standard Error
price_{t-1}	0.922	0.001	0.868	0.002	0.862	0.002
$(\text{price}_{t-1})^2$			0.001	0.000	0.001	0.000
Q_{t-1} (10^-3)					0.008	0.000
d_{Feb}	0.659	0.073	0.645	0.073	0.817	0.073
d_{Mar}	0.560	0.073	0.545	0.072	0.717	0.072
d_{Apr}	-0.108	0.072	-0.115	0.071	0.049	0.071
d_{May}	0.569	0.071	0.547	0.071	0.726	0.071
d_{Jun}	0.497	0.071	0.476	0.071	0.659	0.071
d_{July}	-0.315	0.071	-0.338	0.070	-0.169	0.070
d_{Aug}	1.227	0.070	1.175	0.070	1.345	0.070
d_{Sep}	0.249	0.070	0.221	0.070	0.391	0.070
d_{Oct}	-0.296	0.069	-0.327	0.069	-0.157	0.069
d_{Nov}	-0.062	0.076	-0.078	0.075	0.082	0.076
d_{Dec}	1.299	0.074	1.280	0.073	1.398	0.073
Constant	0.671	0.059	1.438	0.064	1.359	0.064
R^2	0.924		0.925		0.926	
# Observations	58337		58337		58337	

Notes: OLS Regression of prices on lagged prices and quantities for each software title, pooled across all platforms.

Table 11: Porting Cost Estimates, By Genre

	(i) Action			(ii) Family			(iii) Fighting		
	Estimate	95% CI		Estimate	95% CI		Estimate	95% CI	
PS2	0.000	0.000	0.000	0.000	*0.000*	*0.000*	0.000	*0.000*	*0.000*
XBOX	-1.048	*-1.586*	*0.944*	2.631	*1.656*	*7.682*	1.069	*-2.461*	*3.853*
PS2 & XBOX	0.792	*0.245*	*1.380*	0.424	*-0.343*	*0.424*	0.722	*-4.424*	*0.815*
GC	-1.589	*-2.156*	*-0.220*	2.303	*2.003*	*9.342*	0.650	*0.096*	*5.665*
PS2 & GC	*0.083*	*0.049*	*0.429*	0.059	*-0.084*	*1.398*	0.650	*0.011*	*0.924*
XBOX & GC	-0.736	*-1.332*	*0.862*	3.488	*2.746*	*10.930*	0.864	*-2.529*	*4.284*
All 3	0.849	*0.422*	*1.221*	0.641	*-0.187*	*1.401*	1.254	*-2.890*	*2.185*
c_{PS2}	-0.247	*-0.332*	*-0.009*	0.447	*0.353*	*1.381*	0.314	*0.235*	*0.974*
c_{XB}	-0.131	*-0.230*	*-0.033*	-0.074	*-0.074*	*0.085*	0.023	*0.011*	*0.931*
c_{GC}	0.015	*0.009*	*0.038*	-0.008	*-0.272*	*0.023*	0.135	*0.079*	*0.282*
# Titles	*241*			*100*			*77*		

	(iv) Platformer			(v) Racing			(vi) RPG		
	Estimate	95% CI		Estimate	95% CI		Estimate	95% CI	
PS2	0.000	0.000	0.000	0.000	*0.000*	*0.000*	0.000	*0.000*	*0.000*
XBOX	0.218	*-0.291*	*3.613*	0.413	*0.413*	*1.772*	2.784	*-0.992*	*3.955*
PS2 & XBOX	-0.178	*-0.462*	*0.588*	0.901	*0.691*	*2.686*	-0.320	*-0.355*	*0.694*
GC	2.127	*0.535*	*4.537*	-0.155	*-0.297*	*0.420*	4.010	*-0.373*	*4.920*
PS2 & GC	*1.724*	*0.341*	*1.724*	0.231	*-0.002*	*0.318*	0.031	*-0.055*	*0.554*
XBOX & GC	2.567	*1.467*	*22.392*	0.470	*0.470*	*1.899*	3.326	*-0.398*	*4.930*
All 3	1.756	*0.123*	*2.366*	1.344	*1.015*	*3.327*	0.460	*0.460*	*1.317*
c_{PS2}	0.255	*-0.007*	*0.761*	-0.020	*-0.060*	*0.117*	0.817	*0.029*	*1.061*
c_{XB}	0.066	*-0.063*	*0.115*	-0.139	*-0.475*	*-0.131*	0.211	*0.036*	*0.211*
c_{GC}	-0.228	*-0.321*	*-0.004*	0.002	*-0.005*	*0.035*	0.034	*-0.057*	*0.132*
# Titles	*103*			*197*			*173*		

	(vii) Shooter			(viii) Sports			(ix) Other		
	Estimate	95% CI		Estimate	95% CI		Estimate	95% CI	
PS2	0.000	0.000	0.000	0.000	*0.000*	*0.000*	0.000	*0.000*	*0.000*
XBOX	-2.254	*-5.266*	*1.092*	0.114	*-1.379*	*0.219*	-0.431	*-2.034*	*0.219*
PS2 & XBOX	4.761	*-0.010*	*18.284*	-0.043	*-0.085*	*0.291*	0.062	*-0.447*	*1.265*
GC	-7.267	*-14.554*	*-0.133*	-0.025	*-1.401*	*0.106*	-0.384	*-1.367*	*0.533*
PS2 & GC[a]	-	*-.202*	-	-0.181	*-0.959*	*0.101*	2.446	*0.926*	*3.795*
XBOX & GC	-3.239	*-3.240*	*5.030*	-0.043	*-1.404*	*0.146*	-0.190	*-0.190*	*16.922*
All 3	1.470	*-5.931*	*20.358*	1.183	*0.416*	*1.906*	0.248	*0.174*	*3.841*
c_{PS2}	-0.604	*-2.039*	*-0.077*	0.166	*-0.142*	*0.202*	-0.041	*-0.295*	*0.024*
c_{XB}	-0.407	*-2.134*	*0.042*	0.025	*-0.032*	*0.039*	-0.012	*-0.238*	*0.022*
c_{GC}	0.764	*-0.293*	*2.244*	0.076	*0.012*	*0.226*	-0.003	*-0.265*	*0.003*
# Titles	*163*			*308*			*207*		

Notes: Estimates of θ_C used to specify porting costs in (19) (units in \$M), separately estimated by genre. Instruments: constant, title fixed effect α_k^w, inverse of expected quantity sold (M). 95% confidence intervals are constructed taking 80 sample draws from the empirical distribution of the moment inequalities and re-estimating costs.

[a] No titles observed to have chosen this strategy, so no upper bound is identified. During computation of counterfactual regimes, costs are assumed to be arbitrarily high for this particular genre and strategy choice.

Table 12: Dynamic Network Formation Game: Predicted Fit of Model

		Observed Data	Predicted Data		
			Estimate	Conf. Interval	
Installed Base (M)	PS2	30.07	29.40	*27.91*	*29.59*
	XB	13.32	12.86	*12.48*	*13.32*
	GC	9.83	11.02	*10.53*	*11.22*
% Market Shares	PS2	56.50	55.18	*53.92*	*55.54*
	XB	25.03	24.13	*23.57*	*25.67*
	GC	18.47	20.69	*20.15*	*21.46*
# of Titles	PS2	1161	953	*559*	*1031*
	XB	749	943	*787*	*1067*
	GC	487	795	*556*	*884*
# of "Hit" Titles (Sales > 100K)	PS2	578	585	*427*	*613*
	XB	296	293	*272*	*309*
	GC	290	245	*220*	*251*
# of "Hit" Titles (Sales > 1M)	PS2	67	62	*61*	*64*
	XB	9	8	*7*	*9*
	GC	8	11	*10*	*12*
Total Titles Sold (M)	PS2	305.09	297.22	*252.78*	*304.15*
	XB	118.05	115.07	*105.26*	*123.93*
	GC	79.17	103.45	*87.66*	*107.44*

Notes: Predicted data obtained by fixing strategic platform choices for first-party titles, but allowing third-party titles to re-optimize. Porting cost estimates are from table 11. Confidence intervals are computed by redrawing from the estimated porting cost distribution for multiple sets of instruments, and recomputing a new equilibrium.

Table 13: Counterfactual: Banning Vertical Integration and Exclusivity

		Observed Data	(i) CF #1: First Party Titles Estimate	Conf. Interval		(ii) CF #2: No FP Titles Estimate	Conf. Interval	
Installed Base (M)	PS2	30.07	58.02	*57.09*	*58.42*	37.80	*33.85*	*38.16*
	XB	13.32	8.70	*8.59*	*9.08*	10.03	*9.50*	*11.78*
	GC	9.83	10.06	*8.94*	*10.43*	6.04	*5.60*	*6.39*
% Market Shares	PS2	56.50	75.56	*75.17*	*76.25*	70.17	*66.51*	*71.10*
	XB	25.03	11.32	*11.23*	*11.96*	18.62	*17.94*	*22.87*
	GC	18.47	13.10	*11.78*	*13.58*	11.21	*9.75*	*11.82*
Number of Titles	PS2	1161	1175	*756*	*1200*	1008	*448*	*1083*
	XB	749	1169	*984*	*1259*	963	*810*	*1268*
	GC	487	936	*592*	*982*	502	*142*	*666*
# of "Hit" Titles (Sales > 100K)	PS2	578	997	*731*	*1014*	772	*425*	*806*
	XB	296	187	*173*	*200*	172	*143*	*231*
	GC	290	376	*303*	*396*	69	*32*	*76*
# of "Hit" Titles (Sales > 1M)	PS2	67	325	*315*	*329*	155	*127*	*159*
	XB	9	5	*4*	*5*	0	*0*	*2*
	GC	8	32	*19*	*33*	0	*0*	*0*
Total Titles Sold (M)	PS2	305.09	1474.79	*1380.50*	*1491.80*	609.01	*433.77*	*622.27*
	XB	118.05	69.19	*61.42*	*74.12*	60.04	*50.69*	*87.55*
	GC	79.17	178.32	*127.00*	*192.02*	24.50	*10.06*	*28.70*

Notes: Counterfactual results allow all titles to re-optimize and choose the optimal set of platforms. CF #1 assumes first-party titles are still present; CF #2 eliminates them. Estimates are computed using porting cost estimates from table 11. Confidence intervals are computed by redrawing from the estimated porting cost distribution for multiple sets of instruments, and recomputing a new equilibrium.

Table 14: Estimated Parameters of Negative Binomial Regression

	Variable	Parameter Estimate	Standard Error
$\ln(\mu)$	*Constant*	1.423	0.111
	Q_{jt}	0.003	0.000
	Age	-0.025	0.004
	$Age^2(10^{-2})$	0.032	0.004
	d_{Feb}	0.476	0.126
	d_{Mar}	1.109	0.118
	d_{Apr}	0.416	0.127
	d_{May}	0.424	0.128
	d_{Jun}	0.789	0.124
	d_{Jul}	0.322	0.134
	d_{Aug}	0.800	0.125
	d_{Sep}	1.597	0.116
	d_{Oct}	1.492	0.117
	d_{Nov}	1.700	0.113
	d_{Dec}	0.743	0.126
$\ln(\alpha)$	*Constant*	-0.245	0.292
	d_{PS2}	-2.270	0.476
	d_{XBOX}	-2.502	0.595
	d_{GC}	-3.000	0.748
Num. Obs.		458	

Notes: Coefficients are from negative binomial regression of number of titles released at time t for seven consoles, including four fifth-generation consoles. Only coefficients for sixth-generation consoles are reported for $\ln(\alpha)$.

Figure 1: Hardware Quantity and Prices

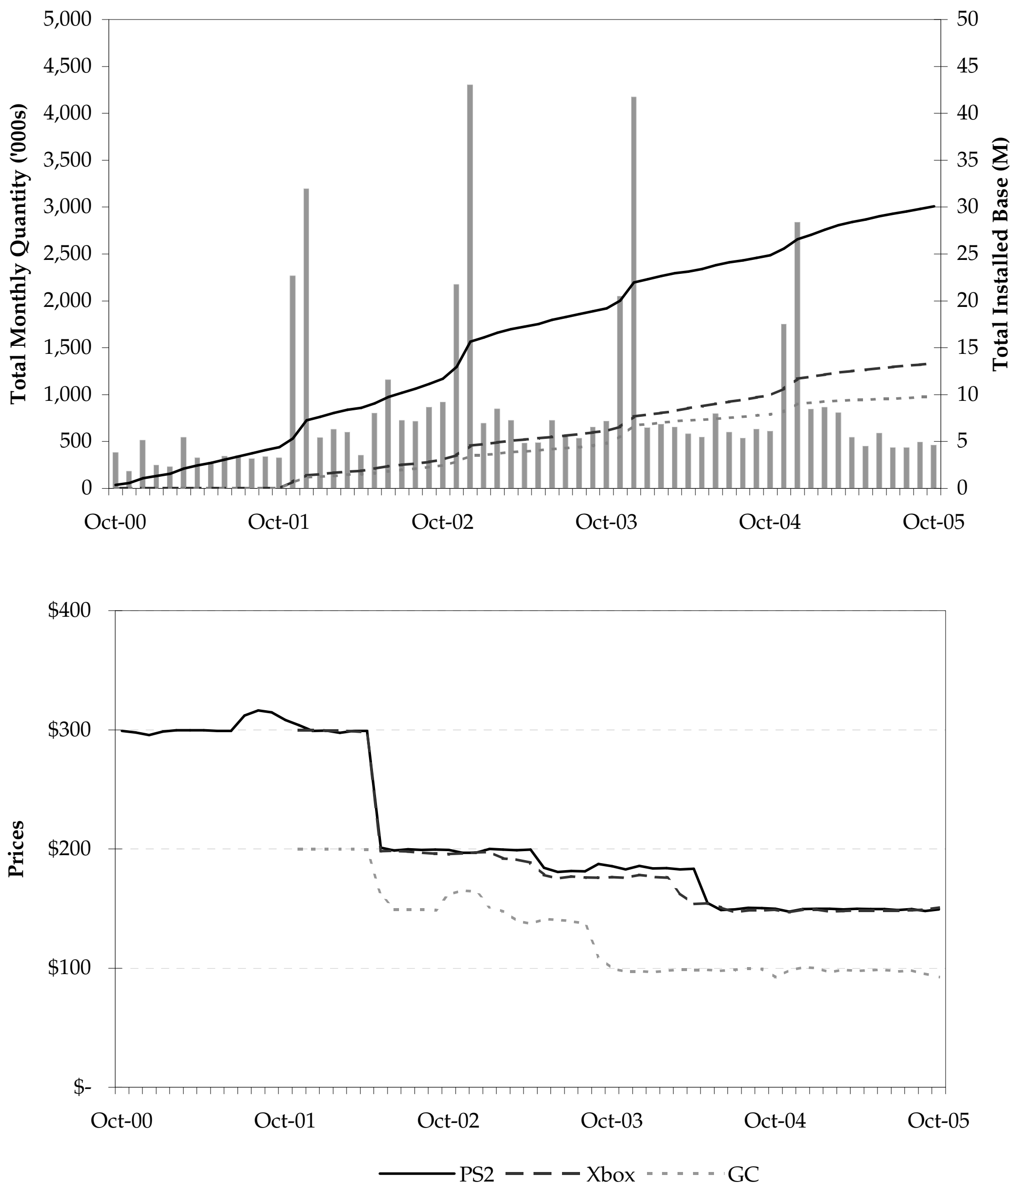

Notes: In the top graph, bars represent the total number hardware consoles sold across all three platforms in each month in thousands (scale on left); lines graphs indicate the total installed base for each console in millions (scale on right). The bottom graph provides average monthly (nominal) prices faced by consumers in retail stores for each platform.

Figure 2: Estimation Algorithm

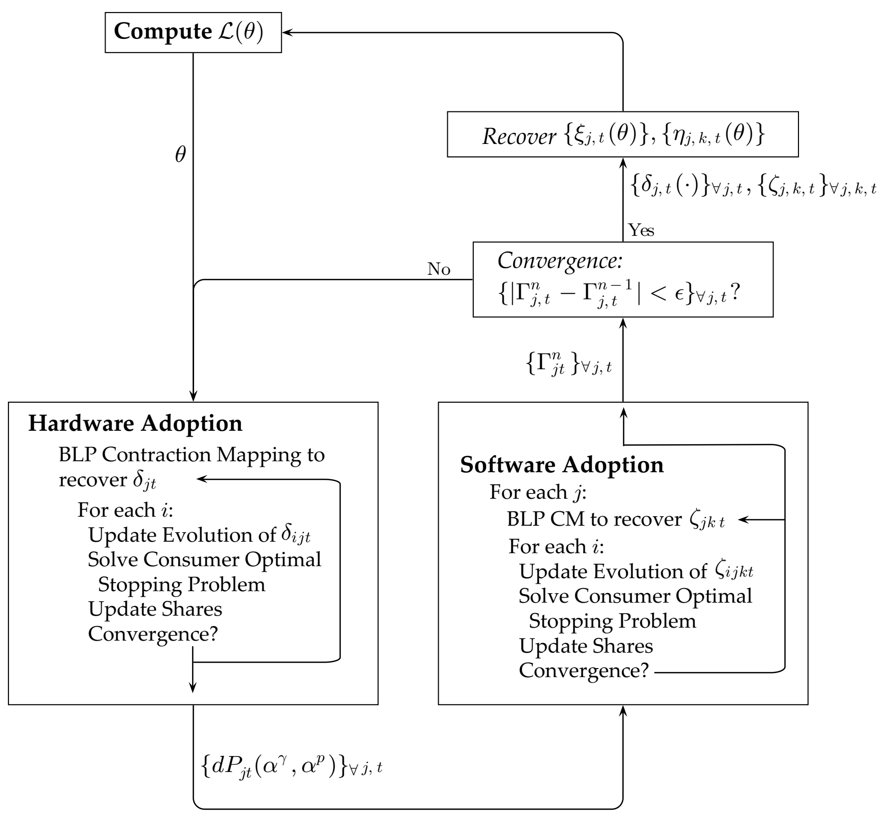

Figure 3: Evolution of Installed Base

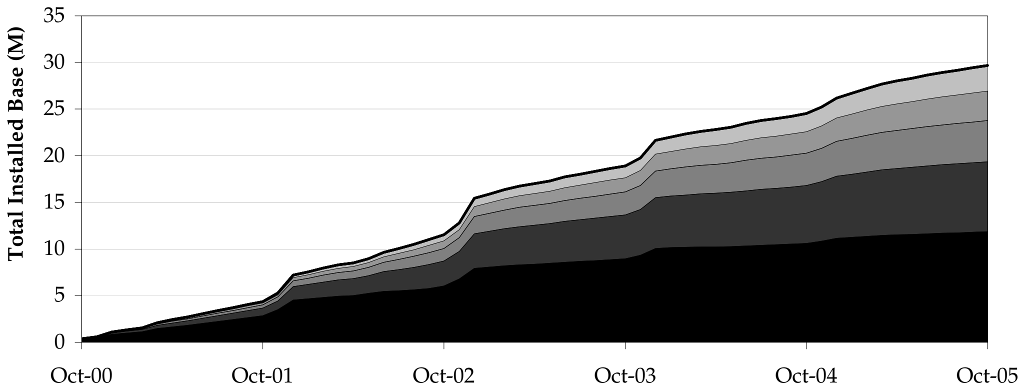

(a) PS2

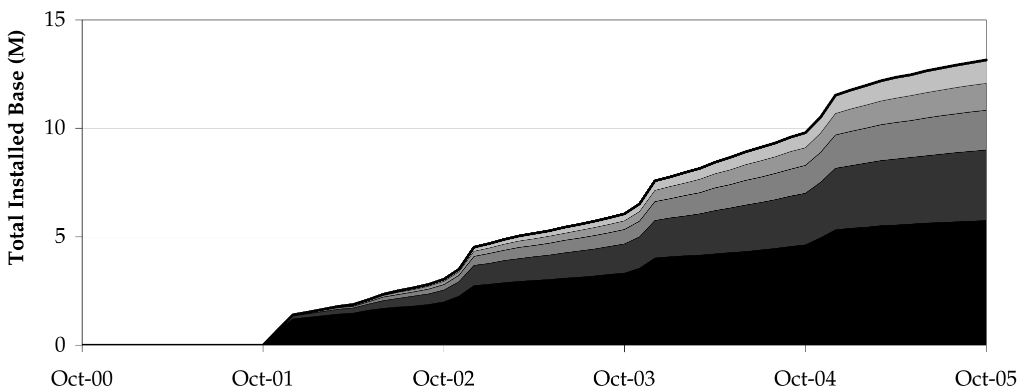

(b) XBOX

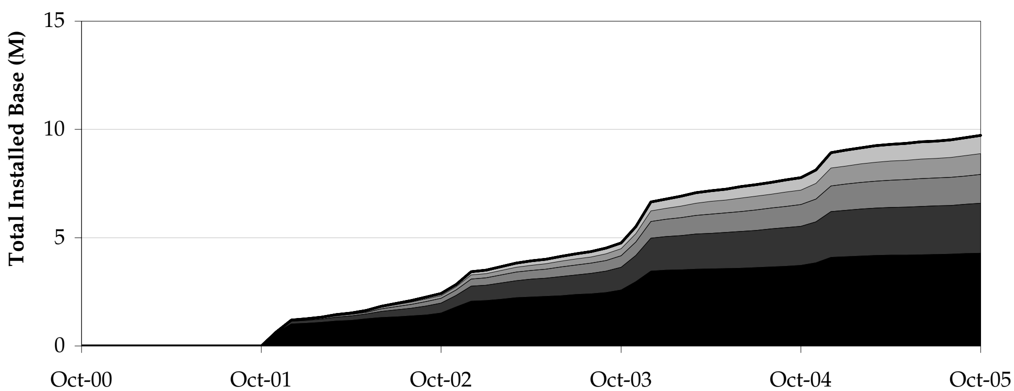

(c) GC

Notes: Implied evolution of installed bases for each console for consumers with different values of α^{γ}. Darkest area at the bottom corresponds to the highest quintile of the distribution; lightest area at top corresponds to lowest quintile.

Figure 4: Estimated Values of $\{\delta_{j,t,0}\}_{\forall j,t}$

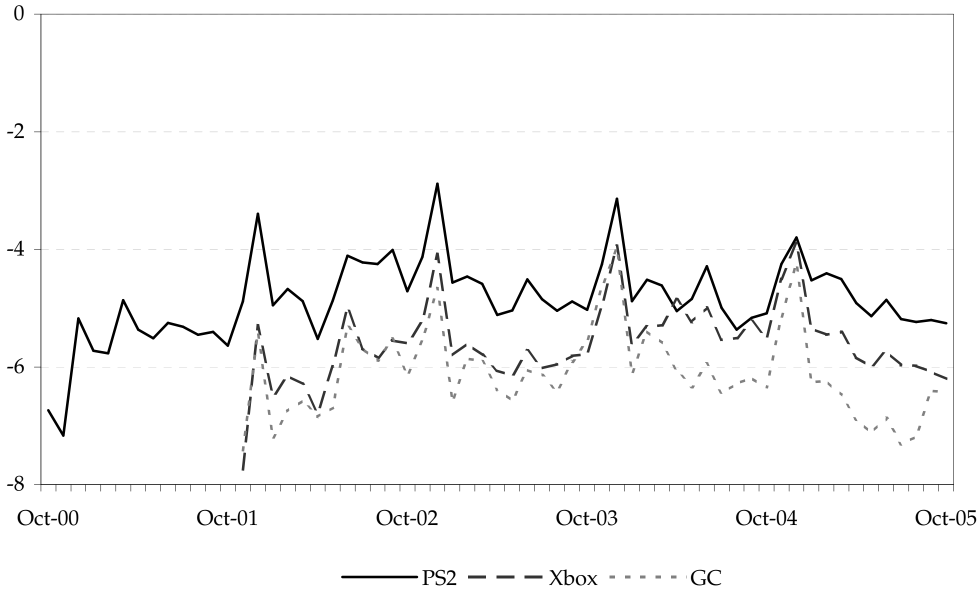

Notes: Realized values of hardware mean-utility δ for mean consumer at inventory state $\iota = 0$ (no previous purchase) implied by the full demand model.

Figure 5: Difference Between Estimated $\delta_{j,t+1,0}$ and Predicted Value $E[\delta_{j,t+1,0}|\{\delta_{j,t,0}\}_{\forall j}, m(t)]$

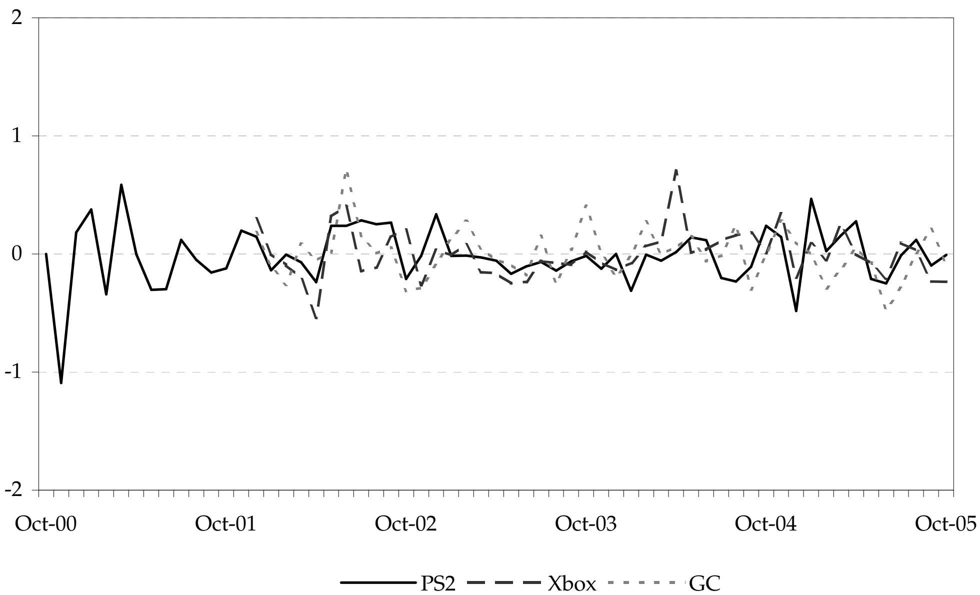

Notes: Errors between realized and predicted values of hardware mean-utility δ for mean consumer with no inventory using the estimated Markov transition process given by (16).

Figure 6: Fit of Model: $\{\nu^{hw}_{j,t}\}_{\forall j,t}$

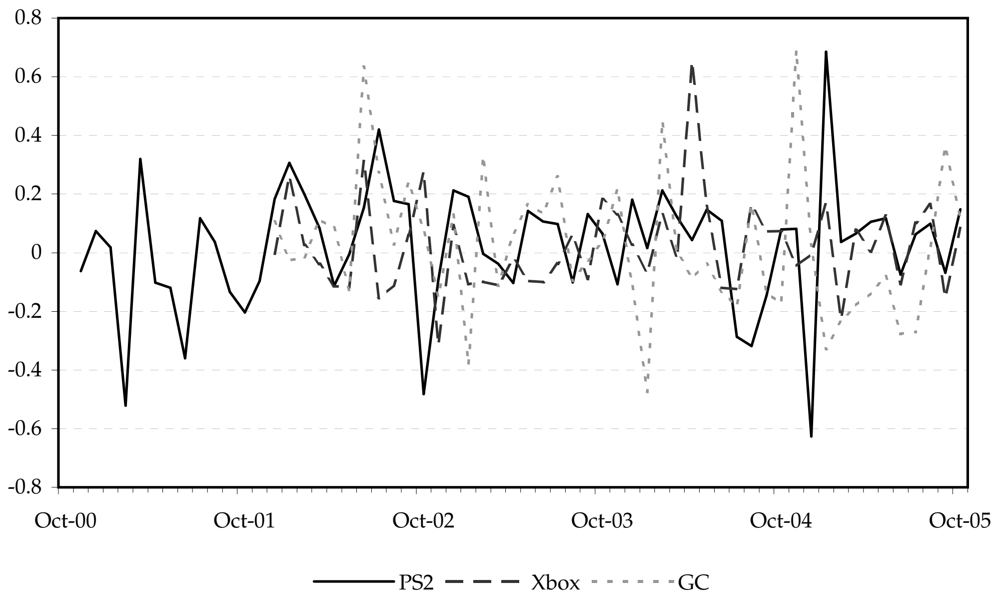

Notes: Predicted residuals in hardware unobserved characteristics from full demand model: $\nu^{hw}_{j,t} \equiv \xi_{j,t} - \rho^{hw}\xi_{j,t}$.

Figure 7: Actual and Predicted Number of Software Titles Released

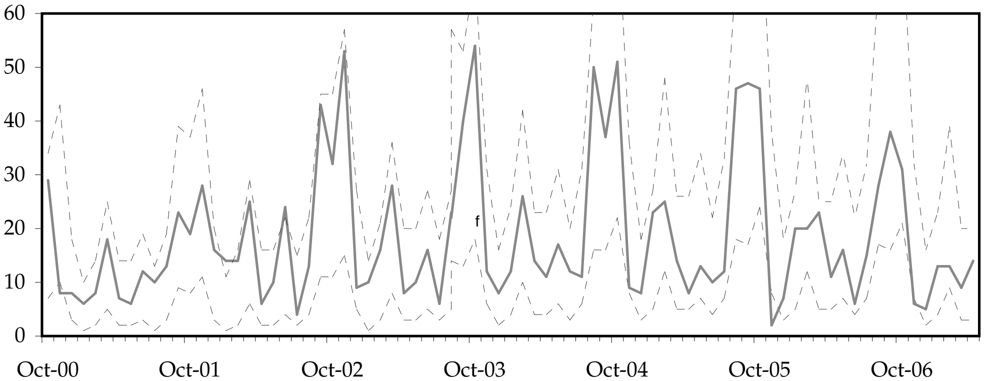

(a) PS2

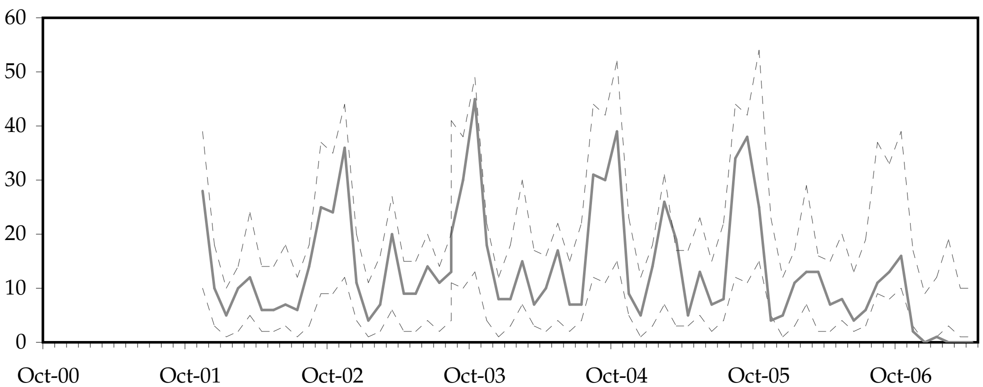

(b) XBOX

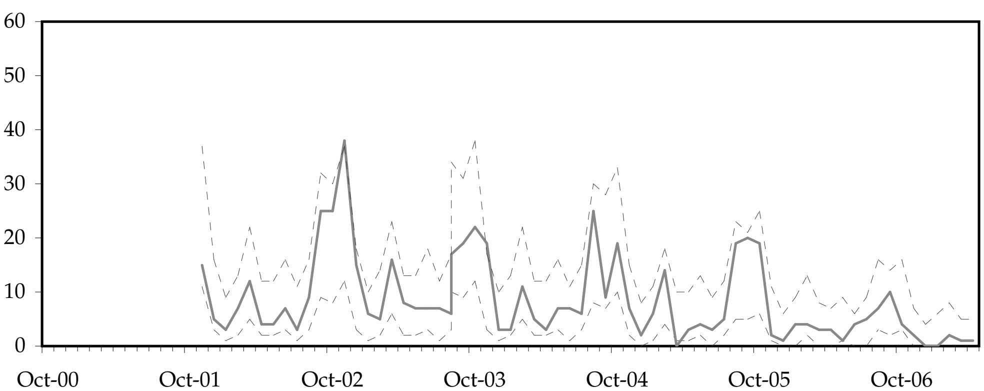

(c) GC

Notes: Solid lines indicate actual number of titles $q_{j,t}$ released for each platform in a given month; dashed lines indicate the 5% and 95% percentiles of the negative binomial distribution with parameters given in table 14 as a function of total software titles available in the previous period $Q_{j,t-q}$.

References

ARMSTRONG, M. (2006): "Competition in Two-Sided Markets," *RAND Journal of Economics*, 37(3), 668–691.

ASKER, J. (2004): "Measuring Advantages from Exclusive Dealing," mimeo.

BEN-AKIVA, M. E. (1973): "The Structure of Travel Demand Models," PhD thesis, MIT.

BENITEZ-SILVA, H., G. HALL, G. J. HITSCH, G. PAULETTO, AND J. RUST (2000): "A Comparison of Discrete and Parametric Approximation Methods for Continuous-State Dynamic Programming Problems," mimeo, Yale University.

BERNHEIM, B. D., AND M. D. WHINSTON (1998): "Exclusive Dealing," *Journal of Political Economy*, 106(1), 64–103.

BERRY, S. (1994): "Estimating Discrete-Choice Models of Product Differentiation," *RAND Journal of Economics*, 25(2), 242–262.

BERRY, S., J. LEVINSOHN, AND A. PAKES (1995): "Automobile Prices in Market Equilibrium," *Econometrica*, 63(4), 841–890.

BERRY, S., O. B. LINTON, AND A. PAKES (2004): "Limit Theorems for Esimating the Parameters of Differentiated Product Demand Systems," *Review of Economic Studies*, 71, 613–654.

BESANKO, D., AND M. K. PERRY (1993): "Equilibrium Incentives for Exclusive Dealing in a Differentiated Products Oligopoly," *RAND Journal of Economics*, 24, 646–667.

BLOCH, F., AND M. JACKSON (2007): "The Formation of Networks with Transfers among Players," *Journal of Economic Theory*, 127(1), 83–110.

BRADLY, R. A., AND J. J. GART (1962): "The Asymptotic Properties of ML Estimators when Sampling from Associated Populations," *Biometrika*, 49(1/2), 205–214.

CAILLAUD, B., AND B. JULLIEN (2003): "Chicken & Egg: Competition Among Intermediation Service Providers," *RAND Journal of Economics*, 34(2), 309–328.

CAMERON, A. C., AND P. K. TRIVEDI (1998): *Regression Analysis of Count Data*. Cambridge University Press, Cambridge, U.K.

CARRANZA, J. E. (2006): "Estimation of Demand for Differentiated Durable Goods," mimeo, University of Wisconsin-Madison.

CHANDA, K. C. (1954): "A Note on the Consistency and Maxima of the Roots of Likelihood Equations," *Biometrika*, 41(1/2), 56–61.

CHIPTY, T. (2001): "Vertical Integration, Market Foreclosure, and Consumer Welfare in the Cable Television Industry," *American Economic Review*, 91(3), 428–453.

CLEMENTS, M. T., AND H. OHASHI (2005): "Indirect Network Effects and the Product Cycle: U.S. Video Games, 1994 - 2002," *Journal of Industrial Economics*, 53(4), 515–542.

CORTS, K. S., AND M. LEDERMAN (2007): "Software Exclusivity and Indirect Network Effects in the US Home Video Game Industry," mimeo, Rotman School of Managment, University of Toronto.

Coughlan, P. J. (2001): "Note on Home Video Game Technology and Industry Structure," Harvard Business School Case 9-700-107.

Dubin, J. A., and D. L. McFadden (1984): "An Econometric Analysis of Residential Electric Appliance Holdings and Consumption," *Econometrica*, 52(2), 345–362.

Einav, L. (2007): "Seasonality in the U.S. Motion Picture Industry," *RAND Journal of Economics*, 38(1).

Ericson, R., and A. Pakes (1995): "Markov-Perfect Industry Dynamics: A Framework for Empirical Work," *Review of Economic Studies*, 62, 53–82.

Evans, D. (2003): "The Antitrust Economics of Multi-Sided Platform Markets," *Yale Journal on Regulation*, 20, 325–382.

Evans, D. S., A. Hagiu, and R. Schmalensee (2006): *Invisible Engines: How Software Platforms Drive Innovation and Transform Industries.* MIT Press, Cambridge, MA.

Farrell, J., and P. Klemperer (2007): "Coordination and Lock-In: Competition with Switching Costs and Network Effects," in *Handbook of Industrial Organization*, ed. by M. Armstrong, and R. Porter, vol. 3. North-Holland Press, Amsterdam.

Farrell, J., and G. Saloner (1985): "Standardization, Compatibility, and Innovation," *RAND Journal of Economics*, 16(1), 70–83.

Fuller, W. A. (1996): *Introduction to Statistical Time Series, 2d Ed.* Wiley, New York.

Gandal, N., M. Kende, and R. Rob (2000): "The Dynamics of Technological Adoption in Hardware/Software Systems: the Case of Compact Disc Players," *RAND Journal of Economics*, 31, 43–61.

Gentzkow, M. (2007): "Valuing New Goods in a Model with Complementarity: Online Newspapers," *American Economic Review*, 97(3).

Goolsbee, A., and A. Petrin (2004): "The Consumer Gains from Direct Broadcast Satellites and the Competition with Cable TV," *Econometrica*, 72, 351–381.

Gowrisankaran, G., and M. Rysman (2007): "Dynamics of Consumer Demand for New Durable Goods," mimeo.

Hagiu, A. (2006): "Pricing and Commitment by Two-Sided Platforms," *RAND Journal of Economics*, 37(3), 720–737.

Hagiu, A., and R. S. Lee (2007): "Exclusivity and Control," mimeo, Harvard Business School. http://www.people.hbs.edu/rlee/papers.html.

Hausman, J., B. H. Hall, and Z. Griliches (1984): "Econometric Models for Count Data with an Application to the Patents-R & D Relationship," *Econometrica*, 52(4), 909–938.

Hayashi, F. (2000): *Econometrics.* Princeton University Press, Princeton, NJ.

Hendel, I., and A. Nevo (2006): "Measuring the Implications of Sales and Consumer Inventory Behavior," *Econometrica*, 74(6), 1637–1673.

Ho, K. (2006): "The Welfare Effects of Restricted Hospital Choice in the US Medical Care Market," *Journal of Applied Econometrics*, 21(7), 1039–1079.

——— (2007): "Insurer-Provider Networks in the Medical Care Market," mimeo.

Ishii, J. (2005): "Interconnection Pricing, Compatibility, and Investment in Network Industries: ATM Networks in the Banking Industry," mimeo.

Katz, M., and C. Shapiro (1986): "Technology Adoption in the Presence of Network Externalities," *Journal of Political Economy*, 94(4), 822–841.

——— (1994): "Systems Competition and Network Effects," *Journal of Economic Perspectives*, 8(2), 93–115.

Kent, S. L. (2001): *The Ultimate History of Video Games*. Three Rivers Press, New York, NY.

Klein, B. (1988): "Vertical Integration as Organizational Ownership: The Fisher Body–General Motors Relationship Revisited," *Journal of Law, Economics and Organization*, 4, 199–213.

Lancaster, K. (1971): *Consumer Demand: A New Approach*. Columbia University Press, New York, NY.

Lee, L. (1995): "Asymptotic Bias in Simulated Maximum Likelihood Estimation of Dsicrete Choice Models," *Econometric Theory*, 11(3), 437–483.

Lee, R. S. (2006): "Competing Platforms," mimeo, Harvard University. http://www.people.hbs.edu/rlee/papers.html.

Magnac, T., and D. Thesmar (2002): "Identifying Dynamic Discrete Decision Processes," *Econometrica*, 70(2), 801–816.

Marvel, H. P. (1982): "Exclusive Dealing," *Journal of Law and Economics*, 25, 1–25.

Mathewson, G. F., and R. A. Winter (1987): "The Competitive Effects of Vertical Agreements," *American Economic Review*, 77, 1057–1062.

McFadden, D. (1973): "Conditional Logit Analysis of Qualitative Choice Behavior," in *Frontiers in Econometrics*, ed. by P. Zarembka, vol. 3. Academic Press, New York, NY.

Melnikov, O. (2001): "Demand for Differentiated Products: The case of the U.S. Computer Printer Market," mimeo.

Nair, H. (2007): "Intertemporal Price Discrimination with Forward-looking Consumers: Application to the US Market for Console Video-Games," *Quantitative Marketing and Economics*, forthcoming.

Nelder, J. A., and R. Mead (1965): "A Simplex Method for Function Minimization," *Computer Journal*, 7, 308–313.

Nevo, A. (2001): "Measuring Market Power in the Ready-to-Eat Cereal Industry," *Econometrica*, 69(2), 307–342.

Pakes, A. (1986): "Patents as Options: Some Estimates of the Value of Holding European Patent Stocks," *Econometrica*, 54(4), 755–784.

Pakes, A., and P. McGuire (1994): "Computing Markov Perfect Nash Equilibrum: Numerical Implications of a Dyamic Differentiated Product Model," *RAND Journal of Economics*, 25(4), 555–589.

Pakes, A., J. Porter, K. Ho, and J. Ishii (2006): "Moment Inequalities and Their Application," mimeo, Harvard University.

Petrin, A. (2002): "Quantifying the Benefits of New Products: The Case of the Minivan," *Journal of Political Economy*, 110(4), 705–729.

Prat, A., and A. Rustichini (2003): "Games Played Through Agents," *Econometrica*, 71(4), 989–1026.

Prieger, J. E., and W.-M. Hu (2006): "An Empirical Analysis of Indirect Network Effects in the Home Video Game Market," NET Institute Working Paper #06-25.

Rao, C. R. (1973): *Linear Statistcial Inference and Its Appliactions: Second Edition*. John Wiley & Sons, New York.

Rasmusen, E. B., J. M. Ramseyer, and J. S. Wiley (1991): "Naked Exclusion," *American Economic Review*, 81(5), 1137–1145.

Rey, P., and J. Tirole (2007): "A Primer on Foreclosure," in *Handbook of Industrial Organization*, ed. by M. Armstrong, and R. Porter, vol. 3. North-Holland Press, Amsterdam.

Rochet, J.-C., and J. Tirole (2002): "Cooperation Among Competitors: Some Economics of Payment Card Associations," *RAND Journal of Economics*, 33(4), 549–570.

——— (2003): "Platform Competition in Two-Sided Markets," *Journal of the European Economic Association*, 1(4), 990–1029.

——— (2006): "Two-Sided Markets: A Progress Report," *RAND Journal of Economics*, 37(3), 645–667.

Rust, J. (1987): "Optimal replacement of GMC bus engines: An empirical model of Harold Zurcher," *Econometrica*, 55, 999–1033.

——— (1994): "Structural Estimation of Markov Decision Processes," in *Handbook of Econometrics*, ed. by R. F. Engle, and D. L. McFadden, vol. 4. North Holland, Amsterdam.

Rysman, M. (2004): "Competition Between Networks: A Study of the Market for Yellow Pages," *Review of Economic Studies*, 71, 483–512.

Saloner, G., and A. Shepard (1995): "Adoption of Technologies With Network Effects: An Empirical Examination of the Adoption of Automated Teller Machines," *RAND Journal of Economics*, 26, 479–501.

Sass, T. (2005): "The Competitive Effects of Exclusive Dealing: Evidence from the U.S. Beer Industry," *International Journal of Industrial Organization*, 23, 203–225.

Segal, I., and M. D. Whinston (2000): "Exclusive Contracts and Protection of Investments," *RAND Journal of Economics*, 31, 603–633.

Shapiro, C. (1999): "Exclusivity in Network Industries," *George Mason Law Review*, 7, 673–683.

Shy, O. (2001): *The Economics of Network Industries*. Cambridge University Press, Cambridge, U.K.

Song, I., and P. K. Chintagunta (2003): "A Micromodel of New Product Adoption with Heterogeneous and Forward Looking Consumers: Application to the Digital Camera Category," *Quantitative Marketing and Economics*, 1(4), 371–407.

Sweeting, A. (2007): "Dynamic Product Repositioning in Differentiated Product Industries: The Case of Format Switching in the Commercial Radio Industry," mimeo.

Takahasi, D. (2002): *Opening the Xbox*. Prima Publishing, Roseville, CA.

Tarone, R. E., and G. Gruenhage (1975): "A Note on the Uniqueness of Roots of the Likelihood Equations for Vector-Valued Parameters," *Journal of the American Statistical Association*, 70(352), 903–904.

Tiebout, C. (1956): "A Pure Theory of Local Expenditures," *Journal of Political Economy*, 64(5), 416–424.

Train, K. E. (2003): *Discrete Choice Methods with Simulation*. Cambridge University Press, Cambridge, U.K.

Wald, A. (1949): "Note on the Consistency of the Maximium Likelihood Estimate," *The Annals of Mathematical Statistics*, 20(4), 595–601.

Weintraub, G. Y., C. L. Benkard, and B. Van Roy (2007): "Markov Perfect Industry Dynamics with Many Firms," mimeo.

Whinston, M. D. (2001): "Exclusivity and Tying in *U.S. V. Microsoft*: What We Know, and Don't Know," *Journal of Economic Perspectives*, 15(2), 63–80.

——— (2006): *Lectures on Antitrust Economics*. The MIT Press, Cambridge, MA.