

Tipping and Concentration in Markets with Indirect Network Effects*

Jean-Pierre Dubé
Graduate School of Business
University of Chicago

Günter J. Hitsch
Graduate School of Business
University of Chicago

Pradeep Chintagunta
Graduate School of Business
University of Chicago

February 25, 2008
First Version: July 1, 2007

Abstract

This paper develops a framework to measure “tipping”—the increase in a firm’s market share dominance caused by indirect network effects. Our measure compares the expected concentration in a market to the hypothetical expected concentration that would arise in the absence of indirect network effects. In practice, this measure requires a model that can predict the counter-factual market concentration under different parameter values capturing the strength of indirect network effects. We build such a model for the case of dynamic standards competition in a market characterized by the classic hardware/software paradigm. To demonstrate its applicability, we calibrate it using demand estimates and other data from the 32/64-bit generation of video game consoles, a canonical example of standards competition with indirect network effects. In our example, we find that indirect network effects can lead to a strong, economically significant increase in market concentration. We also find important roles for beliefs on both the demand side, as consumers tend to pick the product they expect to win the standards war, and on the supply side, as firms engage in penetration pricing to invest in growing their networks.

*We are grateful to Dan Akerberg, Lanier Benkard, Ariel Burstein, Tim Conley, Ron Goettler, Wes Hartmann, Ali Hortaçsu, Harikesh Nair, Amil Petrin, and Peter Rossi for helpful comments and suggestions. We also benefited from the comments of seminar participants at Duke University, Tillburg University, the University of Minnesota, the University of Pennsylvania, the 2007 Invitational Choice Symposium at the Wharton School, the 2007 QME Conference in Chicago, the 2007 SICS Conference at Berkeley and the 2008 winter I.O. meetings at the NBER. We would like to acknowledge the Kilts Center for Marketing and the Initiative for Global Markets (IGM), both at the University of Chicago, for providing research funds. The first author was also supported by the Neubauer Family Faculty Fund. All correspondence may be addressed to the authors at the University of Chicago, Graduate School of Business, 5807 South Woodlawn Avenue, Chicago, IL 60637; or via e-mail at jdube@gsb.uchicago.edu or guenter.hitsch@chicagogsb.edu.

1 Introduction

We study the diffusion of competing durable goods in a market exhibiting indirect network effects due to the classic hardware/software structure (Katz and Shapiro 1985). Of particular interest is whether such markets are prone to *tipping*: “the tendency of one system to pull away from its rivals in popularity once it has gained an initial edge” (Katz and Shapiro 1994) and, in some instances, to emerge as the de facto industry standard. Thus, tipping can create a natural form of market concentration in hardware/software markets. The potential for tipping can also lead to aggressive standards wars between incompatible hardware products as they compete for market dominance. These standards wars are widely regarded as a “fixture of the information age” (Shapiro and Varian 1999).

The extant literature has yet to provide an empirically practical measure of tipping. Therefore, we propose a dynamic framework with which to measure tipping, its relationship to indirect network effects and its ability to lead to market concentration in actual markets. We also use the framework to conduct computational exercises with which to understand the general role of expectations during a standards war, both on the supply and demand sides, and to see how they can push a market to tip in favor of one standard. We expect the analysis of tipping and its natural tendency towards market concentration to be of general importance both to practitioners and to policy makers.¹

The potential for tipping figures prominently in current antitrust discussions about hardware/software markets, as highlighted in the recent high-profile case surrounding the browser war between Microsoft and Netscape (United States v. Microsoft, 87 F. Supp. 2d 30 and Bresnahan 2001). However, existing antitrust policies and tools are often inadequate for addressing the feedback dynamics in markets with indirect network effects (e.g. Evans 2003, Koski and Kretschmer 2004, Evans and Schmalensee 2007, and Rysman 2007). Since adoption decisions are not instantaneous, an empirically relevant model of a hardware/software market needs to incorporate dynamics in demand and supply. These dynamics constitute a methodological challenge and, consequently, much of the extant empirical literature either estimates the effects of indirect network effects using demand only, or treats the supply side of the market as static (Gupta et al. 1999, Basu et al. 2003, Bayus and Shankar 2003, Ohashi 2003, Dranove and Gandal 2003, Nair et al. 2004, Karaca-Mandic 2004, Park 2004, Rysman 2004, Clements and O’Hashi 2005, Akerberg and Gowrisankaran 2007, and Tucker and Ryan 2007). Gandal et al. (2000) allow for forward-looking consumers; but they assume hardware sponsors do not have a strategic role. More recently, Liu (2007) and, most closely-related to

¹Herein, we focus only on the standardization that may emerge from competition. We do not consider the role of formal standard-setting committees (e.g. Katz and Shapiro 1994) such as those that ultimately settled the standards war in the 56K modem market (Augereau et al 2006). We do not consider the role of compatibility which, in some instances, may eliminate the dominance of one of the hardware standards (Chen, Doraszelski and Harrington 2007).

our work, Jenkins, Liu, Matzkin, and McFadden (2004) allow for forward-looking hardware manufacturers. But, both papers treat consumers as myopic. In contrast, our paper allows for forward-looking consumer behavior and solves for an equilibrium in which consumers' and firms' expectations are mutually consistent. We will demonstrate that the assumption of static consumer behavior can strongly limit the empirical relevance of a model for measuring concentration or "bad acts," as forward-looking consumer behavior can strongly exacerbate tipping.

Before discussing our empirical formulation of tipping, we first explain how indirect network effects can lead to tipping. In a hardware/software market structure, indirect network effects arise because consumers adopt hardware based on the current availability and their beliefs about the future availability of software, while the (third-party) supply of software increases in the installed base of a given hardware standard (Chou and Shy 1992, Church and Gandal 1993).² Hence, each standard becomes more valuable to a consumer if it attains a larger installed base. Due to *positive feedback*, a small initial market share advantage can eventually lead to large differences in the shares of the competing standards. This process is exacerbated by rational, self-fulfilling expectations, which allow consumers to coordinate on a standard that is widely adopted based on mutually consistent beliefs about the current and future adoption decisions of other consumers.³ In an extreme case, two standards A and B could be identical ex-ante, but due to self-fulfilling expectations either "all consumers adopt A" or "all consumers adopt B" could be an equilibrium. Hence, due to the emergence of positive feedback and the role of expectations, markets with indirect network effects may become concentrated, i.e. tip towards one of the competing standards.

We now explain why the formulation of an empirically relevant definition of tipping is difficult. Consider first the case of an ex-ante symmetric market where firms face identical demand functions, have the same production costs, etc. In this case, we could measure tipping by comparing the ex-post asymmetry in market shares with the perfectly symmetric outcome where all firms share the market equally.⁴ In actual markets, however, product differentiation, differences in costs, and other differences between standards frequently lead to asymmetric market outcomes, even in the absence of indirect network effects. Hence, we propose a measure of tipping that compares the expected concentration in a market to the hypothetical expected

²Rochet and Tirole (2003) argue that most network effects arise in an indirect manner.

³The role of *coordination* and *expectations* in driving adoption decisions has been a central theme in the theoretical literature on network effects since the seminal work of Katz and Shapiro (1985). For excellent surveys, see Farrell and Klemperer (2006) and Katz and Shapiro (1994).

⁴Note that we focus herein on tipping and market dominance during a specific hardware generation. A related theoretical literature has also studied whether tipping can create inertia across hardware generations when there are innovations (e.g. Farrell and Saloner 1986, Katz and Shapiro 1992 and Markovitch 2004). Therein, tipping, or "excess inertia," is defined by the willingness of consumers to trade-off the scale benefits of a current standard with a large installed base in favor of a new technology without an installed base. Interestingly, in this type of environment, network effects may also serve as a potential barrier to entry (Cabral 2007).

concentration that would arise if the sources of indirect network effects were reduced or eliminated. The key insight is that tipping generally needs to be measured relative to a well-defined, counter-factual market outcome. For an empirical implementation of this measure, we need a model that captures indirect network effects, can be calibrated from actual data, and allows us to make predictions about the equilibrium adoption of the competing standards under various different parameter values capturing the strength of indirect network effects.

To implement the proposed measure of tipping, we build a dynamic model that captures indirect network effects and gives consumer expectations a central role. Our model involves three types of players: consumers, hardware manufacturers, and software developers. The demand side of our model extends the framework of Nair et al. (2004) and allows for dynamic adoption decisions. Consumers are assumed to “single-home,” meaning they adopt at most one of the competing hardware standards.⁵ The utility of each hardware standard increases in the availability and variety of complementary software. Consumers form beliefs about future hardware prices and software availability. These beliefs influence when consumers adopt (the rate of diffusion) and which standard they adopt (the size of each installed base). On the supply side, forward-looking hardware firms compete in prices while anticipating the impact of hardware sales on the future provision of software and, hence, future hardware sales. Software firms provide a variety of titles that is increasing in the installed base of a hardware standard. Our solution concept for this model is Markov perfect Bayesian equilibrium. The complexity of the model makes analytical solution methods intractable, and hence we solve the model numerically.

To demonstrate our model and how it can be used to measure tipping, we calibrate it with demand parameter estimates and other market data from the 32/64-bit generation of video game consoles.⁶ The video game console market is a canonical example of indirect network effects. Furthermore, from previous empirical research, the 32/64-bit generation is known to exhibit indirect network effects (Venkatesh and Bayus 2003, Clements and Ohashi 2005).

Demand estimation per se is not the main point of this paper. But we do need to overcome some econometric challenges to obtain preference estimates that can be used to calibrate our model. These challenges arise from the incorporation of forward-looking consumer behavior. A nested fixed point approach (Rust 1987) would impose a formidable computational burden that is exacerbated by the presence of indirect network effects in the model. Instead, we adapt the two-step procedures of Bajari, Benkard and Levin (2007) (hereafter BBL) and Pesendorfer and Schmidt-Dengler (2006) (hereafter PS-D) to solve our demand estimation problem. A

⁵Recent literature has begun to study the theoretical implications of multi-homing whereby consumers may adopt multiple standards and software firms may create versions for multiple standards (Armstrong 2005).

⁶This approach follows in the tradition of Benkard (2004), Dubé, Hitsch, and Manchanda (2005), and Dubé, Hitsch, and Rossi (2007) by conducting counter-factual simulations of the market outcomes using empirically obtained parameters.

similar approach has recently been employed by Ryan and Tucker (2007).⁷

The calibrated model reveals that the 32/64 bit video game console market can exhibit economically significant tipping effects, given our model assumptions and the estimated parameter values. The market concentration, as measured by the 1-firm concentration ratio in the installed base after 25 periods, increases by at least 23 percentage points due to indirect network effects. We confirm the importance of consumer expectations as an important source of indirect network effects; in particular, we find that tipping occurs at a (monthly) discount factor of 0.9, but not for smaller discount factors. Our model also predicts *penetration pricing* (for small levels of the installed base) if indirect network effects are sufficiently strong. In markets with strong network effects, firms literally price below cost during the initial periods of the diffusion to invest in network growth.

2 Model

We consider a market with competing hardware platforms. A consumer who has adopted one of the available technologies derives utility from the available software for that platform. Software titles are incompatible across platforms. Consumers are assumed to choose at most one of the competing hardware platforms and to purchase software compatible with the chosen hardware, a behavior Rochet and Tirole (2003) term “single-homing.” There are indirect network effects in this market, which are due to the dependence of the number of available software titles for a given platform on that platform’s installed base. The consumers in this market have expectations about the evolution of hardware prices and the future availability of software when making their adoption decisions. Correspondingly, the hardware manufacturers anticipate the consumer’s adoption decisions, and set prices for their platforms accordingly. The software market is monopolistically competitive, and the supply of software titles for any given platform is increasing in the platform’s installed base.

Time is discrete, $t = 0, 1, \dots$. The market is populated by a mass $M = 1$ of consumers. There are $J = 2$ competing firms, each offering one distinct hardware platform. $y_{jt} \in [0, 1]$ denotes the installed base of platform j in period t , i.e., the fraction of consumers who have adopted j in any period previous to t . $y_t = (y_{1t}, y_{2t})$ describes the state of the market.

In each period, platform-specific demand shocks ξ_{jt} are realized. ξ_{jt} is private information to firm j , i.e., firm j learns the value of ξ_{jt} before setting its price, but learns the demand shock of its competitor only once sales are realized. As we shall see later, ξ_{jt} can strongly influence the final distribution of shares in the installed base. In particular, the realizations of ξ_{jt} in the

⁷An interesting difference is that Ryan and Tucker (2007) use individual level adoption data, which enables them to accommodate a richer treatment of ‘observed’ consumer heterogeneity. The trade-off from incorporating more heterogeneity is that they are unable to solve the corresponding dynamic hardware pricing game on the supply side.

initial periods of competition can lead the market to “tip” in favor of one standard. Also, ξ_{jt} will typically ensure that the best response of each firm is unique, and thus the existence of a pure strategy equilibrium.⁸ We assume that the demand shocks are independent and i.i.d. through time. $\phi_j(\cdot)$ denotes the pdf of ξ_j , and $\phi(\cdot)$ denotes the pdf of $\xi = (\xi_1, \xi_2)$.

The timing of the game is as follows:

1. Firms learn their demand shock ξ_{jt} and set a product price, p_{jt} .
2. Consumers adopt one of the available platforms or delay their purchase decisions.
3. For each platform j , software firms supply a given number of titles, n_{jt} .
4. Sales are realized, and firms receive their profits. Consumers derive utility from the available software titles and—in the case of new adopters—from the chosen platform.

Software Market

The number of available software titles for platform j in each period is a function of the installed base of platform j : $n_{jt} = h_j(y_{j,t+1})$. To see why n_{jt} is a function of $y_{j,t+1}$ and not y_{jt} , note that y_{jt} denotes the installed base at the beginning of period t , while $y_{j,t+1}$ denotes the total installed base after the potential adopters have made a purchase decision. The software producers observe this total installed base before they supply a given number of titles.

While this derivation may seem ad hoc, in Appendix A we show how this relationship between n_{jt} and $y_{j,t+1}$ can be derived from a structural model of monopolistic competition and CES software demand in the software market. These assumptions abstract away from some of the dynamic aspects of game demand (e.g. Nair 2006), but they retain the fundamental inter-dependence between software and hardware.

Consumer Decisions

Consumers make their adoption decisions based on current prices and their expectation of future prices and the availability of compatible software titles. Consumers expect that the installed hardware evolves according to $y_{t+1} = f^e(y_t, \xi_t)$, and that firms set prices according to the policy function $p_{jt} = \sigma_j^e(y_t, \xi_{jt})$. Consumers observe both ξ_t and the current price vector p_t before making their decisions.

Consumer who have already adopted one of the platforms receive utility from the available software in each period. As the supply of software is a function of the installed base at the

⁸We are not able to prove this statement in general, but could easily verify it across all versions of our model that we solved on a computer. In general, the right hand side of the firm’s Bellman equation, regarded as a function of p_{jt} , has two local maxima. The realization of ξ_{jt} ensures that these local maxima are not equal.

end of a period, we can denote this utility as $u_j(y_{j,t+1}) = \gamma n_{jt} = \gamma h_j(y_{j,t+1})$. The present discounted software value is then defined as

$$\omega_j(y_{t+1}) = \mathbb{E} \left[\sum_{k=0}^{\infty} \beta^k u_j(y_{j,t+1+k}) | y_{t+1} \right].$$

This value follows the recursion

$$\omega_j(y_{t+1}) = u_j(y_{j,t+1}) + \beta \int \omega_j(f^e(y_{t+1}, \xi)) \phi(\xi) d\xi.$$

Consumers who have not yet adopted either buy one of the hardware platforms or delay adoption. The choice-specific value of adopting hardware platform j is given by

$$v_j(y_t, \xi_t, p_t) = \delta_j + \omega_j(f^e(y_t, \xi_t)) - \alpha p_{jt} + \xi_{jt}. \quad (1)$$

Here, δ_j is the value of owning a specific hardware platform, or the value of bundled software. α is the marginal utility of income. The realized utility from adopting j also includes a random utility component ϵ_{jt} , which introduces horizontal product differentiation among the competing standards. That is, the total utility from the choice of j is given by $v_j(y_t, \xi_t, p_t) + \epsilon_{jt}$. We assume that ϵ_j is i.i.d. Type I Extreme Value distributed.

The value of waiting is given by

$$v_0(y_t, \xi_t) = \beta \int \max \left\{ v_0(y_{t+1}, \xi) + \epsilon_0, \max_j \{ v_j(y_{t+1}, \xi, \sigma^e(y_{t+1}, \xi)) + \epsilon_j \} \right\} \phi_\epsilon(\epsilon) \phi(\xi) d(\epsilon, \xi). \quad (2)$$

In this equation, $y_{t+1} = f^e(y_t, \xi_t)$.

Consumers choose the option that yields the highest choice-specific value, including ϵ_{jt} . That is, option j is chosen if and only if for all $k \neq j$, $v_j(y_t, \xi_t, p_t) + \epsilon_{jt} \geq v_k(y_t, \xi_t, p_t) + \epsilon_{kt}$.⁹ Given the distributional assumption on the random utility component, the market share of option j is

$$s_j(y_t, \xi_t, p_t) = \frac{\exp(v_j(y_t, \xi_t, p_t))}{\exp(v_0(y_t, \xi_t)) + \sum_{k=1}^J \exp(v_k(y_t, \xi_t, p_t))}. \quad (3)$$

Furthermore, the installed base of platform j evolves according to

$$y_{j,t+1} = y_{jt} + \left(1 - \sum_{k=1}^J y_{kt} \right) s_j(y_t, \xi_t, p_t) = f_j(y_t, \xi_t, p_t). \quad (4)$$

⁹These inequalities involve some slight abuse of notation, as $v_0(y, \xi)$ is not a function of p .

Firms

Firms set prices according to the Markovian strategies $p_j = \sigma_j(y, \xi_j)$, i.e., prices depend only on the current payoff-relevant information observed to each competitor. Firms expect that the consumers make adoption decisions according to the value functions v_0, \dots, v_J , and accordingly that market shares are realized according to equation (3) and that the installed base evolves according to (4).

The marginal cost of hardware production is c_j , which we assume to be constant through time. The firms also collect royalty fees from the software manufacturers at the rate of r_j per unit of software. Let $q_j(y_{t+1})$ be the total number of software titles sold in period t .¹⁰ The per-period expected profit function is then given by

$$\begin{aligned} \pi_j(y, \xi_j, p_j) &= (p_j - c_j) \cdot \left(1 - \sum_{k=1}^J y_{kt}\right) \int s_j(y, \xi_j, \xi_{-j}, p_j, \sigma_{-j}(y, \xi_{-j})) \phi_j(\xi_{-j}) d\xi_{-j} \\ &\quad + r_j \int q_j(f_j(y, \xi_j, \xi_{-j}, p_j, \sigma_{-j}(y, \xi_{-j}))) \phi_j(\xi_{-j}) d\xi_{-j}. \end{aligned}$$

Each competitor maximizes the expected present discounted value of profits. Associated with the solution of the inter-temporal pricing problem is the Bellman equation

$$V_j(y, \xi_j) = \sup_{p_j \geq 0} \left\{ \pi_j(y, \xi_j, p_j) + \beta \int V_j(f(y, \xi_j, \xi_{-j}, p_j, \sigma_{-j}(y, \xi_{-j})), \xi'_j) \phi(\xi_{-j}) \phi(\xi'_j) d(\xi_{-j}, \xi'_j) \right\}. \quad (5)$$

Equilibrium

We seek a Markov perfect Bayesian equilibrium, where firms and consumers base their decision only on the current payoff-relevant information. Consumers have expectations about future hardware prices and the evolution of the installed base of platform and the associated supply of software. The adoption decisions are dependent on these expectations. Firms have expectations about the adoption decisions of the consumers, the evolution of the installed base, and the pricing decisions of their competitors. Pricing decisions are made accordingly. In equilibrium, these expectations need to be mutually consistent.

Formally, a *Markov perfect Bayesian equilibrium* in pure strategies of the network game consists of consumer expectations f^e and σ^e , consumer value functions v_k , pricing policies σ_j , and the firm's value function V_j such that:

1. The consumer's choice-specific value functions v_1, \dots, v_J satisfy (1), and the value of waiting, v_0 , satisfies (2).

¹⁰ $q_j(y_{t+1}) = \bar{q}_j(n_{jt}) \cdot y_{j,t+1}$, where \bar{q}_j denotes the average number of titles bought by a consumer. See Appendix A for the derivation of this equation within the context of a specific model. For our model simulations, we estimate $q_j(y)$ directly from the data.

2. The firm's value functions V_1, \dots, V_J satisfy the Bellman equations (5).
3. $p_j = \sigma_j(y, \xi_j)$ maximizes the right-hand side of the Bellman equation (5) for each $j = 1, \dots, J$.
4. The consumer's expectations are rational: $\sigma_j^e \equiv \sigma_j$ for $j = 1, \dots, J$, and $f^e(y, \xi) = f(y, \xi, \sigma(y, \xi))$, where f is as defined by equation (4).

In the Markov perfect equilibrium, all players—firms and consumers—act rationally given their expectations about the strategies of the other market participants. Furthermore, expectations and actually realized actions are consistent.

3 Estimation

To make our computational results more realistic, we calibrate them with data from the video game console market. While demand estimation per se is not the main objective of the paper, it is nevertheless helpful to discuss briefly some of the challenges involved in estimating preference parameters for a dynamic discrete choice model. The main difficulty arises from the incorporation of consumer beliefs, a crucial element for durable goods demand in general (Horsky 1990, Melnikov 2000, Song and Chintagunta 2003, Nair 2005, Prince 2005, Carranza 2006, Gowrisankaran and Rysman 2006, and Gordon 2006). Once we include consumer beliefs in our console demand function, the derivation of the market shares, equation (3), requires us to compute the choice-specific value functions. Nesting the corresponding dynamic programming problem into the estimation problem is prohibitive due to the high dimension of the state space (y_t, ξ_t and other exogenous states included in the empirical specification). In addition, the derivation of the density of market shares (or moments of the density) requires inverting the demand shocks, ξ , out of the market share function numerically. This step is also computationally costly since ξ enters the utility function non-linearly through the value functions $\omega_j(f^e(y_t, \xi_t))$ and $v_0(y_t, \xi_t)$.¹¹

Instead, we follow a recent tradition in the empirical literature on dynamic games and estimate the structural parameters of our model in two stages (e.g. BBL, PS-D, and Aguirregabira and Mira 2002, 2006). The goal is to construct moment conditions that match the observed consumer choices in the data with those predicted by our model. Rather than computing the choice-specific value functions needed to evaluate demand, we instead devise a two-step approach to simulate them.

¹¹ As discussed in Gowrisankaran and Rysman (2006), since we do not know the shape of these value functions a priori, it is unclear whether the market share function is invertible in ξ , let alone whether a computationally fast contraction-mapping can be used to compute the inverse.

Stage 1

In the first stage, we estimate the consumer choice strategies along with the firms' pricing strategies and the software supply function. The supply function of software variety is specified as follows

$$\log(n_{jt}) = \mathcal{H}_j(y_{j,t+1}; \theta_n) + \eta_{jt}. \quad (6)$$

where $\eta_{jt} \sim N(0, \sigma_\eta^2)$ captures random measurement. This specification is consistent with the equilibrium software supply function derived in the Appendix, equation (9). The pricing strategies are specified as follows

$$\log(p_{jt}) = \mathcal{P}_j(y_t, z_t^p; \theta_p) + \lambda \xi_{jt}, \quad (7)$$

where $\xi_{jt} \sim N(0, 1)$. In equation (7) we let \mathcal{P}_j be a flexible functional form of the state variables. For the empirical model, we include exogenous state variables, z_t^p , that are observed by console firms in addition to y_t and ξ_t , the state variables in the model of Section 2. These additional states are discussed in Section 4. In equation (7), we assume that the video game console manufacturers use only payoff-relevant information to set their prices. But we do not assume that their pricing strategies are necessarily optimal. This specification has the advantage that it is consistent with the Bayesian Markov perfect equilibrium concept used in our model, but does not explicitly impose it.

Conditional on the model parameters, there is a deterministic relationship between the price and installed base data and the demand unobservable, $\xi_{jt} = \mathcal{X}_j(y_{jt}, p_{jt}, z_t^p)^{12}$. Then, conditional on y_t and p_t , we can estimate the consumers' optimal choice strategy in log-odds:

$$\begin{aligned} \mu_{jt} &\equiv \log(s_{jt}) - \log(s_{0t}) \\ &= v_j(y_t, \xi_t, p_t, z_t^d) - v_0(y_t, \xi_t, z_t^d) + \zeta_{jt} \\ &= \mathcal{L}_j(y_t, \mathcal{X}(y_t, p_t, z_t^p), z_t^d; \theta_\mu) + \zeta_{jt}, \end{aligned} \quad (8)$$

where $\zeta_{jt} \sim N(0, \sigma_\zeta^2)$ is random measurement error and z_t^d denotes exogenous state variables observed by the consumer. By including the *control function* $\mathcal{X}(y_t, p_t, z_t^p)$ in the demand equation, we also resolve any potential endogeneity bias that would arise due to the correlation between prices and demand shocks (this is the control function approach used in . The first stage consists then of estimating the vector of parameters $\Theta = (\theta_n, \theta_p, \theta_\mu, \lambda)$ via maximum likelihood using the equations (6), (7), and (8).

¹²We can trivially invert ξ out of the price equation because of the additivity assumption in (7). This is a stronger condition than in BBL, but it is analogous to other previous work such as Villas-Boas and Winer (1999) and Petrin and Train (2005).

Stage 2

In the second stage, we estimate the consumers' structural taste parameters, Λ , by constructing a minimum distance procedure that matches the simulated optimal choice rule for the consumers to the observed choices in the data. The idea is to use the estimated consumer choice strategies, (8) and the laws of motion for prices and software variety, (7) and (6), to forward-simulate the consumers' choice-specific value functions, $\mathcal{V}_j(y_t, \xi_t, p_t; \Lambda, \hat{\Theta})$ and $\mathcal{V}_0(y, \xi; \Lambda, \hat{\Theta})$. The details for the forward-simulation are provided in the Appendix. Note that while our two-step approach does not require us to assume that firms play the Markov Perfect equilibrium strategies explicitly, we do need to assume that consumers maximize the net present value of their utilities.

The minimum distance procedure forces the following moment condition to hold approximately:

$$Q_{jt}(\Lambda_0, \hat{\Theta}) \equiv \mu_{jt} - \left(\mathcal{V}_j(y, \xi, p; \Lambda_0, \hat{\Theta}) - \mathcal{V}_0(y, \xi; \Lambda_0, \hat{\Theta}) \right) = 0.$$

That is, at the true parameter values, Λ_0 , and given a consistent estimate of Θ , the simulated log-odds ratios should be approximately equal to the observed log-odds ratios for each of the observed states in the data. The minimum distance estimator, Λ^{MD} , is obtained by solving the following minimization problem:

$$\Lambda^{MD} = \min_{\Lambda} \left\{ Q(\Lambda, \hat{\Theta})' W Q(\Lambda, \hat{\Theta}) \right\},$$

where W is a positive semi-definite weight matrix.¹³ Wooldridge (2002) shows that the minimum distance estimator has an asymptotically normal distribution with the covariance matrix

$$Avar(\Lambda^{MD}) = (\nabla_{\Lambda} Q' W \nabla_{\Lambda} Q)^{-1} \nabla_{\Lambda} Q' W \nabla_{\Theta} Q \hat{\Omega} \nabla_{\Theta} Q' W \nabla_{\Lambda} Q (\nabla_{\Lambda} Q' W \nabla_{\Lambda} Q)^{-1},$$

where $\hat{\Omega} = Avar(\hat{\Theta})$, and $\nabla_{\Lambda} Q$ and $\nabla_{\Theta} Q$ denote gradients of Q with respect to Λ and Θ respectively.

The approach is closest to PS-D. But, our implementation differs in two ways. First, we examine a model with continuous states (PS-D look at a model with discrete states). Second, we adapt the approach to estimation of aggregate dynamic discrete choice demand, whereas PS-D focus on discrete choice at the individual level.

¹³We just set W equal to the identity matrix since it is unclear how to derive the efficient W in closed form for our specific problem.

4 Data

For our calibration, we use data from the 32/64-bit generation of video game consoles, one of the canonical examples of indirect network effects. To understand the relevance of this case study to our model and our more general point about tipping in two-sided markets, we briefly outline some of the institutional details of the industry. We then discuss the data.

The US Videogame Console Industry

The market for home video game systems has exhibited a two-sided structure since the launch of Atari's popular 2600 VCS console in 1977 (Williams 2002). Much like the systems today, the VCS consisted of a console capable of playing multiple games, each on interchangeable cartridges. While Atari initially developed its own proprietary games, ultimately more than 100 independent developers produced games for Atari and more than 1,000 games were released for Atari 2600 VCS (Coughlan 2001A). This same two-sided market structure has characterized all subsequent console generations, including the 32/64-bit generation we study herein.

The 32/64-bit generation was novel in several ways. None of the consoles were backward-compatible, eliminating concerns about a previously-existing installed base of consumers which might have given a firm an advantage. This was also the first generation to adopt CD-ROM technology; although early entrants, Philips and 3DO, failed due to their high console prices of \$1000 and \$700 respectively. In contrast, the September 1995 US launch of Sony's 32-bit CD-ROM console, Playstation, was an instant success. So much so, that its competitors, Sega's 32-bit Saturn console and later, Nintendo's 64-bit N64 cartridge console, failed to recapture Sony's lead. In fact, Sega's early exit from the market implied a duopoly console market between Sony's first-generation Playstation and Nintendo's N64.

Playstation's success reflected several changes in the management of the console side of the market. From the start, Sony's strategy was to supply as many games as possible, a lesson it learned from its experience with Betamax video technology:

Sony's primary goal with respect to Playstation was to maximize the number and variety of games... Sony was willing to license any Playstation software that didn't cause the hardware to "crash." — Coughlan (2001b)

To stimulate independent game development, Sony charged substantially lower game royalties of \$9, in contrast with Nintendo's \$18 (Coughlan 2001c). Sony's CD-based platform also lowered game development costs, in contrast with Nintendo's cartridge based system. While the Playstation console failed to produce any truly blockbuster games during its first year (Kirkpatrick 1996), after three months, Playstation's games outnumbered those of Sega's

Saturn by three-to-one. By 1998, more than 400 Playstation titles were available in the US. In addition, Sony engaged in aggressive penetration pricing of the console early on, hoping to make its money back on game royalties (Cobb 2003).

In contrast, Nintendo maintained very stringent conditions over its game licensees, a legacy from its management of game licensees during earlier generations when Nintendo was dominant.¹⁴ By Christmas of 1996, N64 only had eight games in contrast with roughly 200 Playstation titles (Rigdon 1996). By June 1997, N64 still had only 17 games while Playstation had 285. Nintendo insisted that it competed on quality, rather than quantity and in 1997 its CEO claimed, “Sony could kill off the industry with all of ‘its garbage’” (Kunii 1998). In the end, the dominance of Sony Playstation in the 32/64-bit console generation was attributed primarily to its vast library of games, rather than to specific game content.

Recall that our case study focuses only on the 32/64-bit console generation. The success of Playstation’s game proliferation strategy makes us comfortable with the assumption that game variety proxies meaningfully for the indirect network effects. This assumption would be more tenuous for more recent console generations now that blockbuster games have become more substantial. For example, the blockbuster game Halo 3, for Microsoft’s Xbox 360, generated \$300 million in sales during its first week (Blakely 2007) and, by November 2007, it represented over 17% of Xbox 360’s worldwide game sales according to NPD. At the same time, monthly Xbox 360 console sales nearly doubled in contrast with two months previously, selling 527, 800 units in October 2007 (Gallagher 2007). Similarly, Playstation 3’s Spiderman 3 grossed \$151 million during its first week (Blakely 2007). The blockbuster games of the 32/64-bit generation were smaller in magnitude. Only three N64 games garnered over 4% of total US game unit sales on the N64 platform, Goldeneye 007, Mario Kart 64 and Super Mario 64, while an additional 21 games captured over 1% of total game sales. Only five Playstation titles captured over 1% of total Playstation game sales, none capturing over 2%.¹⁵ Nair (2007) tests for Blockbuster game effects during this generation. He finds no material impact on sales or prices of games in the months leading-up to the launch of a best-selling game. Therefore, Nair (2007) ignores competitive effects in his analysis of video game pricing during this generation.

Data

Our data are obtained from NPD Techworld’s “Point of Sale” database. The database consists of a monthly report of total sales and average prices for each video game consoles across a

¹⁴The dominance of Nintendo’s 8-bit NES console, during the 1980s, allowed it to command 20% royalties in addition to a manufacturing fee of \$14 per game cartridge. Licensees were also restricted to 5 new NES titles per year. Nevertheless, by 1991, less than 10% of titles were produced by Nintendo and the system had over 450 titles in the US. In addition, one in three US households had an NES console by 1991, with the average console owner purchasing 8 or 9 games (Coughlan 2001B).

¹⁵These numbers are based on US game sales data collected by NPD.

sample of participating US retailers from September 1995 to September 2002. NPD states that the sheer size of the participating firms represent about 84% of the US retail market. We also observe the monthly number of game titles available during the same period. We define the potential market size as the 97 million US households as reported by the US Census.

In the data, we observe a steady decline in console prices over time. At first glance, this pattern seems inconsistent with the penetration-pricing motive one would expect from our model. However, Playstation is estimated to have launched at a price roughly \$40 below marginal cost (Coughlan 2001) and console prices have been documented to have fallen more slowly than costs over time, the latter due to falling costs of chips (Liu 2006). The rising margins over time are consistent with penetration pricing. Although we do not observe marginal costs, we control for falling costs by including a time trend as a state in the empirical model. Thus, our empirical model is consistent with a richer game in which firms face falling marginal costs. We include this time trend in both z_t^p and z_t^d , which treats it as a commonly-observed state. In addition, we experiment with producer price indices (PPI's), from the BLS, for computers, computer storage devices and audio/video equipment to control for technology costs associated with a console. Finally, we also experiment with the inclusion of the exchange rate (Japanese Yen per US dollar) to control for the fact that parts of the console are sourced from Japan. These two sets of cost-shifting variables, PPI's and exchange rates, are included in z_t^p . However, since we do not expect these costs to be observed by consumers when they make console purchase decisions, we exclude them from z_t^d .

The empirical model also includes monthly fixed-effects to control for the fact that there are peak periods in console demand (e.g. around Christmas). These states are observed by both firms and consumers and, hence, enter both z_t^p and z_t^d . For the policy simulations, we will ignore the effects of time, month and cost-shifters since they are incidental to our theoretical interest in tipping.

Descriptive statistics of the data are provided in Table 1. The descriptive statistics indicate a striking fact about competition between Sony Playstation and Nintendo 64. On average, the two consoles charged roughly the same prices. However, Sony outsold Nintendo by almost 50%. At the same time, over 3.5 times as many software titles were available for Sony than for Nintendo. Of interest is whether Sony's share advantage can be attributed to its large pool of software titles.

Identification

Like most of the extant literature estimating structural models of durable goods demand, our diffusion data contain only a single time-series for the US market¹⁶. The use of a single time-

¹⁶An interesting exception is Gupta et al (1999), who use panel data on individual HDTV adoption choices obtained from a conjoint experiment.

series creates several generic identification concerns for durable goods demand estimation in general¹⁷. The first and most critical concern is the potential for sales diffusion data to exhibit dependence over time as well as inter-dependence in the outcome variables. In addition, the diffusion implies that any given state is observed at most once, a property that could complicate the estimation of beliefs. Finally, we also face the usual potential for price endogeneity to bias demand parameters if prices are correlated with the demand shocks, ξ (Berry 1994). We now briefly discuss the intuition of our empirical identification strategy.

Diffusion data may naturally exhibit dependence over time in prices, p_t , and an inter-dependence between prices and the other outcome variables, y_t and n_t . A concern is whether we can separately identify the price coefficient, α , and the software taste (i.e. the indirect network effect), γ . Our solution consists of adding console cost-shifting variables, PPI's and exchange rates, that vary prices but that are excluded from demand and from software supply. The exclusion restrictions introduce independent variation in prices and, hence, in the term αp_t in the utility function. The exchange rates are particularly helpful in this regard because they introduce independent variation over time — past research has documented that short-run exchange rate innovations follow a random walk (e.g. Meese and Rogoff 1982 and Rogoff 2007). The exclusion restrictions embody a plausible assumption that consumers do not observe the PPI's and exchange rates and, hence, they do not adjust their expectations in response to them.

A related concern is whether we can separately identify the role of product differentiation, δ_j (i.e. one standard has a higher share due to its superior technology), and the indirect network effects, γ (i.e. one standard has a higher share due to its larger installed base which in turn stimulates more software variety) on demand. Our assumption of “single-homing” (i.e. discrete choice), a reasonable assumption for this generation of video game consoles, enables us to infer preferences from aggregate market shares. In addition, we hold each console's quality fixed over time. Thus, we can identify the current utility of software (i.e. the indirect network effect) using variation in the beginning-of-period installed base, y_t .

Finally, we face the usual concerns about endogeneity bias due to prices (e.g. Berry 1994). We do not have a specific console attribute or macro taste shock in mind when we include ξ in the specification; but we include it as a precautionary measure. We are reasonably confident ξ is not capturing the impact of unmeasured blockbuster games¹⁸. Nevertheless, to the extent that ξ captures demand information that is observed by firms, any resulting correlation

¹⁷Some argue that data containing multiple independent markets resolves some of the identification issues. Pooling markets would certainly resolve some of our identification concerns; however, it also raises others. Pooling markets requires the strong assumptions that all markets are in the same long-run equilibrium and that all markets have the same parameters (e.g. consumer tastes are the same across markets) in order to estimate beliefs.

¹⁸We checked the correlation between the ξ estimates from our first stage and the 1-firm concentration ratio of video game sales for each console (based on NPD data). Game concentration explains less than 1% of the variation in Playstation's ξ , versus 11% of N64's ξ .

between prices and ξ could introduce endogeneity bias. Our joint-likelihood approach to the first stage does provide a parametric solution to the endogeneity problem through functional form assumptions. We have imposed a structure on the joint-distribution of the data which provides us with the relationship between prices and demand shocks, ξ . However, we can relax this strong parametric condition by using our console cost-shifters. Both the exchange rate and the PPI's provide sources of exogenous variation in prices that are excluded from demand and that are unlikely to be correlated with consumer tastes for video game consoles, i.e. ξ . In essence, the endogeneity is resolved by including the control function, $X(y_{jt}, p_{jt}, z_t^p)$, in the log-odds of choices, equation (8) (e.g. Villas-Boas and Winer 1999 and Petrin and Train 2007).

5 Estimation Results

5.1 First Stage

During the first stage, we experiment with several specifications. In Table 2, we report the log-likelihood and Bayesian information criterion (BIC) associated with each specification.

Our findings indicate that allowing the states to enter h_j both linearly and quadratically improves fit substantially based on the BIC predictive fit criterion (model 3 versus model 2). Allowing for a time trend also improves fit moderately (model 2 versus model 1). We use a time-trend that is truncated after 65 periods since prices roughly level off after that point (i.e. we do not expect costs to decline indefinitely). We also experimented with a more flexible distributional assumption for the demand shocks, ξ . We use a mixture-of-normals specification to check whether the assumption of normality potentially biases our MLE's. However, we find little change in fit from the 2-component mixture (model 4 versus model 3). Moving to the last three rows, models 5, 6 and 7, we look at the implications of including additional cost proxies into the pricing function that are excluded from the game supply and from the consumer choices. Recall these are terms we include in z_t^p , but we do not include in z_t^d . We use a 3-month lag and 7-month lag in the exchange rate as they were found to explain more price variation than the contemporaneous exchange rate, which is likely due to the fact that production is sourced in advance of sales. Overall, the inclusion of these terms in the price equation improves the overall likelihood of the first stage (as seen by the BIC for model 7).

Although not reported in the Tables, a regression of log-prices on the various price-shifters, including the PPI's and the exchange rate, generates an R^2 of 0.9. Similarly, the OLS regression for the game titles generates an R^2 of 0.98. In the case of log-odds, the inclusion of ξ makes it hard to interpret an R^2 . Instead, we construct a distribution of ξ using a parametric bootstrap from the asymptotic distribution of the parameters in the price regressions.

The mean R^2 of a regression of log-odds on the observed states and the simulated ξ is 0.95. Overall, the first-stage model appears to fit the data well.

A critical aspect of the 2-step method is that the first-stage model captures the relationship between the outcome variables and the state variables. To assess the fit of the first-stage estimates, in Tables 3, 4, and 5, we report all the first-stage estimates and their standard errors. Most of the estimates are found to be significant at the 95% level. In Table 5, we find a positive relationship between software variety and the installed base of each standard. Analogous findings are reported in Clements and O’Hashi (2005).

In the Figures 1, 2, and 3, we plot the true prices, log-odds and games under each standard. In each case, we plot the outcome variable for a standard against its own installed base (reported as a fraction of the total potential market, $M = 97,000,000$). In addition, we report a 95% prediction interval for each outcome variables based on a parametric bootstrap from the asymptotic distribution of our parameter estimates¹⁹. In several instances, the observed outcome variable lies slightly outside the prediction interval. But, overall, our first-stage estimates appear to do a reasonably good job preserving the relationship between the outcome variables and the installed base.

5.2 Second Stage

We report the structural parameters from the second-stage in Table 6. Results are reported for two specifications: models 3 and 7 from the previous section. Recall that model 3 does not have any exclusion restrictions across equations in the first stage. Model 7, the best-fitting model overall in stage 1, includes PPI’s and exchange rates in the price equations. To estimate the second stage of the model, we maintain the assumption that consumers do not observe realizations of these costs. Instead, we assume they observe prices each period and can integrate the innovations to prices out of their expected value functions²⁰. The results are based on an assumed consumer discount factor of $\beta = 0.9$ and 60 simulated histories²¹ of length 500 periods each. Although not reported, we also included monthly fixed-effects in tastes.

First, both model specifications each appear to yield qualitatively similar results. While the point estimates suggest a slight preference for the Sony PlayStation console, the difference in tastes between the two consoles is statistically insignificant. This finding is consistent with

¹⁹The prediction intervals are constructed as follows. 5000 draws are generated from the asymptotic distribution of the first-stage parameter estimates. We then compute the predicted log-price, log-odds and log of game titles corresponding to each parameter draw. We then plot the 5th and 95th percentiles of these values.

²⁰To estimate the distributions of these various costs, we assume they all follow a random walk distribution with drift. Thus, we regress each cost on its 1-period lag along with an intercept and an i.i.d. shock. For the PPI’s, we obtain an R^2 of 0.99, whereas for the exchange rates, we obtain an R^2 of 0.89.

²¹Since the second-stage estimator is linear in the simulation error, the choice of the number of draws only influences efficiency.

industry observers who noted that the improvements from 32 to 64 bit technology were much less dramatic than in previous generations (Coughlan 2002). Rather, the variety of availability of games tended to be the main differentiator. Indeed, the taste for software variety, γ , is positive and significant. In both specifications, γ is roughly 0.1. The effective “network effect” in the model arises from the positive (and significant) software taste on the demand side, γ , and the positive (and significant) elasticity of each standard’s supply of software titles with respect to its installed base, λ_{Sony} and $\lambda_{Nintendo}$ (as in Table 5). The qualitative implications of these estimates are best understood in the context of our simulations in the following section.

6 Model Predictions

We now return to the main questions addressed in this paper regarding the ability to quantify “tipping” using our empirical estimates from the 32/64-bit video game console market. Throughout this section, we will base our simulations on Model 7, which was the best-fitting specification in stage 1. We define tipping as the extent to which the economic mechanism of indirect network effects leads to market concentration. Indirect network effects in our model arise both through the consumers’ current marginal utility of complementary software, γ , and through their discount factor, β , which determines how the consumers’ expectations about the future availability of software for each standard impact on current adoption decisions. Using our empirical example, we examine how γ and β lead to market concentration and tipping.

Our model abstracts from certain aspects of the 32/64 hardware market, in particular learning-by-doing (declining production costs) and persistent heterogeneity in consumer tastes. In this respect, we caution that our predictions should not be interpreted as attempts to explain literally the observed, historic evolution of the market.

In the special case of a market with two symmetric competitors, we can define a measure of tipping by comparing the expected installed base share of the larger standard T periods after the product launch to a market share of 50% (i.e. to the the share in a “symmetric” outcome). That is, we measure tipping as the extent to which the cumulative 1-firm concentration ratio in period T exceeds 50%. In most actual markets, however, the expected share of the larger standard will exceed 50% even in the absence of indirect network effects, due to product differentiation, cost differences across the standards, etc. To assess tipping, we need to compare the expected share in the installed base to the hypothetical share that would arise if one or more economic factors that cause indirect network effects were absent or smaller in size. Therefore, we need a model to predict counter-factual market outcomes, and define the (counter-factual) baseline case relative to which tipping is measured.²²

²²Note that due to demand shocks, the expected cumulative 1-firm concentration ratio could significantly exceed 50% in a market with symmetric competitors even if there are no indirect network effects. In this

We now provide a formal definition of our tipping measure. Let ρ_{jt} be the share of standard j in the installed base t periods after product launch:

$$\rho_{jt} \equiv \frac{y_{j,t+1}}{y_{1,t+1} + y_{2,t+1}}.$$

Here, remember that $y_{j,t+1}$ is the installed base of standard j at the end of period t and thus includes the sales of j during period t . The cumulative 1-firm concentration ratio after T periods is then given by

$$\mathcal{C}(y_T) = \max\{\rho_{1T}, \rho_{2T}\}.$$

The realization of $\mathcal{C}(y_T)$ depends on the model parameters, Θ , an equilibrium that exists for these parameters, $\mathcal{E}(\Theta)$, and a sequence of demand shocks, ξ_t . Given Θ and $\mathcal{E}(\Theta)$, the distribution of $(y_t)_{t=0}^T$ is well defined, and we can thus calculate the expected cumulative 1-firm concentration ratio

$$C_1(\Theta, \mathcal{E}(\Theta)) \equiv \mathbb{E}(\mathcal{C}(y_T) | \Theta, \mathcal{E}(\Theta)).$$

Let Θ' be a variation of the model where one or more parameters that govern the strength of indirect network effects are changed compared to the model described by Θ , and let $\mathcal{E}(\Theta')$ be a corresponding equilibrium. We can thus measure tipping, the increase in market concentration due to indirect network effects, as

$$\Delta C_1 = C_1(\Theta, \mathcal{E}(\Theta)) - C_1(\Theta', \mathcal{E}(\Theta')).$$

If we knew that the market under investigation was symmetric, then $C_1 \approx 0.5$ in the absence of indirect network effects, and we could measure tipping by $\Delta C_1 = C_1(\Theta, \mathcal{E}(\Theta)) - 0.5$.

To implement the tipping measure, we calibrate the model developed in Section 2 and use it to predict the evolution of the market. The parameters consist of the demand estimates and software supply function estimates presented in section 5, along with industry estimates of hardware console production costs and royalty fees.²³ For a given set of parameter values, we solve for a Markov perfect Bayesian equilibrium of the model, and then simulate the resulting equilibrium price and adoption paths.

situation, the difference between the cumulative 1-firm concentration ratio and 50% would not provide a meaningful measure of tipping even in the symmetric case.

²³Cost and royalty data are reported in Liu (2007) and based on various industry reports. The marginal production costs are \$147 (Sony) and \$122 (Nintendo), and correspond to Liu's cost estimates 20 months after the launch of Nintendo 64. The royalty fees per game sold are \$9 (Sony) and \$18 (Nintendo).

Preliminaries

We first summarize specific aspects of the model solutions and simulations. Firms and consumers make decisions at the monthly level. Throughout, we assume that firms discount future profits using the factor $\beta = 0.99$ ²⁴. However, we will consider various consumer discount factors across the different simulations. To simplify the analysis, we also normalize the market size to $M = 1$ ²⁵.

We summarize the firms' equilibrium pricing strategies by the expected pricing policies $\mathbb{E}(p_{jt}|y_t) = \mathbb{E}_{\xi_j}(\sigma_j(y_t, \xi_j)|y_t)$. Here, the expectation is taken over the firm's private information, the transitory demand component ξ_j . The equilibrium evolution of the state vector is summarized by a vector field, where each state is associated with the expected state in the next period.²⁶ Thus, for a given current state y_t , we calculate (and plot) a vector describing the expected movement of the state between periods:

$$\vec{\zeta}_t = \mathbb{E}(y_{t+1}|y_t) - y_t = \mathbb{E}_{\xi} (f(y_t, \xi, \sigma(y_t, \xi))|y_t) - y_t.$$

Using the equilibrium policies and equilibrium state transitions, we can simulate a path of prices, sales, and installed base values given an initial condition y_0 and a sequence of demand shocks, ξ_t . For each set of parameter values, we generate 5,000 simulations of the evolution of the market. Using the simulated values, we can then examine the distribution of prices over time, and the distribution of shares in the total installed base at the end of each period, ρ_{jt} .

Measuring Tipping: Symmetric Competition

We first analyze a case of symmetric competition, where both competitors have identical demand functions, production costs, and royalty fee structures. In the symmetric case, it is easy to compare the predicted market concentration relative to the benchmark case, where both competitors share the market equally. We assume that both competitors are characterized by the parameter estimates that we obtained for Sony. As there are no ex ante differences left between Sony and Nintendo apart from the name, we refer to the two competitors as "Standard 1" and "Standard 2."

We first examine how market outcomes are influenced by the consumers' marginal utility of software, γ . We use the parameter estimates obtained for the consumer discount factor $\beta = 0.9$ and then scale the estimated software utility coefficient by the factors 0.25, 0.5, 0.75, and 1. Figure 4 displays the resulting equilibrium pricing policies and predicted expected price paths for the different software utility values. The expected price paths are conditional on cases where Standard 1 sells at least as many consoles as Standard 2 by the end of period

²⁴This discount factor corresponds to an annual interest rate of 12.8%.

²⁵Note that this normalization also requires re-scaling the parameters in the supply equation accordingly.

²⁶As before, the expectation is taken over the demand shock ξ .

$T = 25$, $y_{T1} \geq y_{T2}$. The marginal production costs are indicated by horizontal lines. Figure 5 shows the vector field describing the expected evolution of the state, and the distribution of shares in the installed base, ρ_{jt} , $T = 25$ months after both standards were launched.²⁷

For the scale factors 0.25, 0.5, and 0.75, the results are similar. Prices rise over time, as firms compete more aggressively when they have not yet obtained a substantial share of the market. After 25 months, both firms have an approximately equal share of all adopters. Hence, market outcomes are approximately symmetric.

Now compare these results to the model solution obtained for the estimated software utility coefficient (scale factor equals 1), indicating a larger indirect network effect than in the previous three model variations. Now, the equilibrium changes not only quantitatively but also qualitatively. First, unlike in the previous cases, we are no longer able to find a symmetric equilibrium in pure strategies. However, there are at least two asymmetric pure strategy equilibria. The graphs at the bottom of Figure 5 display one of these equilibria, which “favors” Standard 1. In this equilibrium, before any consoles have been sold ($y_0 = (0, 0)$), consumers expect that Standard 1 will obtain a larger market share than Standard 2 (note the direction of the arrow at the origin). These expectations are self-fulfilling, and due to the impact of the expected future value of software on adoption decisions, Standard 1 will, on average, achieve a larger share of the installed base than Standard 2. If, on the other hand, Standard 2 ever obtains a share of the installed base that is sufficiently larger than the share of Standard 1 (due to a sequence of favorable demand shocks, for example), then consumers’ expectations flip and Standard 2 is expected to win. The advantage due to self-fulfilling expectations is increasing in the difference of shares in the installed base, $y_{jt} - y_{-j,t}$.

As a consequence of this equilibrium behavior, the market becomes concentrated, even though the standards are identical *ex ante*. The expected cumulative one-firm concentration ratio increases from $C_1 = 0.502$ for the scale factor 0.25 to $C_1 = 0.833$ for the scale factor 1 (see Table 7). The distribution of shares in the installed base not only becomes disperse, but also asymmetric: in about 55% of all simulations, Standard 1 “wins” the market, i.e. obtains a larger share of the installed base than Standard 2. Note that there is also another asymmetric equilibrium which “favors” Standard 2. This equilibrium exactly mirrors the one which favors Standard 1; for example, Standard 2 has a 55% chance of “winning” the market, etc.

Another interesting aspect of the equilibrium is the impact of the magnitude of the marginal utility of software on firms’ pricing strategies. As can be seen at the bottom of Figure 4, for a scale factor of 1, pricing becomes substantially more aggressive than under the smaller scale factors. For small values of y_{jt} , the firms engage in penetration pricing whereby prices are set below costs (the per-console production cost in the simulations is \$147).

Next, we examine how market outcomes change under different values of the consumers’

²⁷The red bar at the abscissa value ρ depicts the percent of all model simulations, i.e., approximately the probability that Standard 1 accounts for a fraction ρ of all adopters at the end of $T = 25$.

discount factor, β . The discount factor influences how consumers value software that they expect to become available in the future, and thus determines the importance of expectations in driving adoption decisions. We choose several discount factors ($\beta = 0.6, 0.7, 0.8, 0.9$) and solve the model for each β , holding the other parameters that were estimated for the discount factor $\beta = 0.9$ constant. Figure 6 shows that the equilibria obtained and the expected concentration of the market is highly sensitive to the magnitude of β . For the smaller discount factors ($\beta < 0.9$), corresponding to relatively small indirect network effects, we obtain a symmetric equilibrium where the expected one-firm concentration ratio C_1 is just slightly larger than 0.5 (Table 7). For $\beta = 0.9$, however, we are unable to compute a symmetric equilibrium, and the expected market concentration increases to $C_1 = 0.833$, as already discussed above.

Alternatively, we derive comparative static result for the same discount factors ($\beta = 0.6, 0.7, 0.8, 0.9$), but solve the model at the parameter values that were estimated for each corresponding β . The results from this exercise are informative on how sensitive the predicted market outcomes are to the choice of the consumers' discount factor, a parameter that is typically *assumed* and not estimated in applied work. The results are similar to the previous ones where we only varied β , but not the other model parameters. In particular, Table 7 shows that the market outcome is almost symmetric for the smaller discount factors, and then becomes very concentrated for $\beta = 0.9$.

Measuring Tipping: The General Case

The symmetric case discussed in the previous section establishes the intuition for the model predictions. We now turn to the measurement of tipping due to indirect network effects in the general case where firms are asymmetric ex ante. With heterogeneous competitors, markets can obviously become concentrated even if indirect network effects are entirely absent. Hence, we measure tipping relative to a specific, counter-factual outcome where one or all mechanisms leading to indirect network effects are absent or small in size.

We first focus on the consumers' software utility parameter, γ . As before in the symmetric case, we scale this parameter by the factors 0.25, 0.5, 0.75, and 1.²⁸ Figure 7 shows the expected market evolution and distribution of shares in the installed base for the different scale factors. Unlike in the case of symmetric competition, one standard, Nintendo, has a persistent advantage for all of the smaller scale factors (0.25, 0.5, and 0.75). In all 5,000 model simulations, Nintendo obtains a larger installed base share than Sony by the end of period $T = 25$, and the expected one-firm concentration ratio, C_1 , ranges from 0.552 and 0.594 (Table 7). At the estimated parameter values (scale factor = 1), however, we once again see a big qualitative and quantitative change in the equilibrium. First, the market becomes

²⁸The consumer discount factor is set to $\beta = 0.9$.

significantly more concentrated, $C_1 = 0.827$. Second, Sony is now predicted to obtain a larger installed base share than Nintendo in 84% of all cases. That is, indirect network effects strongly increase the concentration of the market, and furthermore, the identity of the larger standard changes. The reason for this difference in outcomes for different magnitudes of the indirect network effect is that, according to our estimates, Sony dominates Nintendo in terms of the quantity of software titles supplied at any given value of the installed base. On the other hand, Nintendo has a lower console production cost (\$122 versus \$147). For small values of the software utility, Nintendo’s cost advantage results in lower equilibrium prices and thus a market share advantage over Sony. Once the software utility gives rise to sufficiently large network effects, however, Sony’s advantage in the supply of games becomes important and helps it to “win” the standards war against Nintendo. The same argument also explains why initially, for the scale factors 0.25, 0.50, and 0.75, the concentration ratio C_1 slightly decreases: Sony obtains a larger market share as its relative advantage due to indirect network effects becomes more pronounced.

Next, we examine the market outcomes under different consumer discount factors ($\beta = 0.6, 0.7, 0.8, 0.9$). First, we vary β , but hold all other parameters constant at their estimated values, which were obtained for a discount factor of 0.9. The results (Table 7 and Figure 8) show that the market concentration increases from $C_1 \approx 0.59$ for $\beta < 0.9$ to $C_1 = 0.827$ for $\beta = 0.9$. Furthermore, while—as we already discussed—Sony has a larger share of the installed base than Nintendo in 84% of all cases when $\beta = 0.9$, Nintendo is always predicted to “win” for the smaller discount factors. These predictions remain qualitatively and quantitatively similar when we re-estimate all parameters for each separate consumer discount factor.

Finally, Table 8 shows our measure of tipping, ΔC_1 , for the model predictions at the estimated parameter values relative to several counter-factual models characterized by lower values of the software utility parameter, γ , or smaller values of the consumers’ discount factor, β . For example, compared to a market where the consumers’ flow utility from software is only 25% of the estimated value, indirect network effects are predicted to increase the market concentration by 23 percentage points. Furthermore, relative to a market where consumers discount the future using $\beta = 0.6$, the increase in the market concentration is between 23 and 26 percentage points, depending on the exact counter-factual chosen. Hence, for this particular market, we predict a large, quantitatively significant degree of tipping.

7 Conclusions

We provide a framework for studying the dynamics of hardware/software markets. The framework enables us to construct an empirically practical definition of tipping: the level of concentration relative to a counter-factual in which indirect network effects are reduced or eliminated. Computational results using this framework also provide several important

insights into tipping. Using the demand parameters from the video game industry, we find that consumer expectations play an important role for tipping. In particular, tipping emerges as we strengthen the indirect network either by increasing the utility from software or by increasing the degree of consumer patience. In some instances, this can lead to an increase in market concentration by 23 percentage points or more. Interestingly, tipping is not a necessary outcome of forward-looking behavior. For discount factors as high as 0.8, we observe market concentration falling to roughly the level that would emerge in the absence of any indirect network effects.

Studying other aspects of the equilibrium sheds some interesting managerial insights into the pricing and diffusion. In particular, strengthening the indirect network effect toughens price competition early on during the diffusion, leading firms to engage in penetration-pricing (pricing below marginal cost) to invest in the growth of their networks. When tipping arises, the market diffuses relatively quickly. Thus, an interesting finding is that increasing consumer patience to the point of tipping leads to a more rapid diffusion of consoles.

Our approach to measuring tipping and its role as a source of market concentration should be of interest to antitrust economists, academics and practitioners. For policy workers, our counter-factual approach provides an important method for assessing damages to “bad acts” in markets with indirect network effects. Our results relating consumer and firm beliefs and patience to tipping should also be of interest to academics studying dynamic oligopoly outcomes in markets with durable goods, in particular with indirect network effects. Finally, for the modeling framework constitutes a state-of-the-art quantitative paradigm for practitioners to assess the long-run market share of new durable goods, in particular those exhibiting network effects.

Our main goal herein is to study the role of consumer beliefs and expectations for tipping, not to explain the empirical diffusion of video game consoles per se. Therefore, even though we calibrate the model with data from the 32/64-bit video game console market, we abstract from certain aspects of the industry. For instance, we do not account for declining production costs and persistent consumer heterogeneity when we simulate the market outcomes. Therefore, we caution that our model predictions should not be seen as an attempt to “explain” directly the historical market outcome in the 32/64 bit video game console industry. Nevertheless, studying learning-by-doing, on the supply side, and consumer segmentation, on the demand side, are two interesting directions for future research in this area.

Another area for future research is the role of the game content for console adoption. We intentionally chose the 32/64-bit generation of consoles to allow us to work with a simpler model of the game side of the market. However, during subsequent generations, blockbuster games have become crucial for console adoption decisions. A very interesting direction for future research would be to extend the framework we provide herein to study the role of market power and dynamics on the software side of the model. Similarly, as more recent

generations of game consoles become increasingly targeted (e.g. Nintendo Wii appeals to families while Xbox 360 appeals more narrowly to adult males), households may begin to purchase multiple consoles. Thus, multi-homing may also be an interesting future extension of our framework.

References

- [1] Augereau, Angélique, Shane Greenstein and Marc Rysman (2006), "Coordination versus Differentiation in a Standards War: 56K Modems," *The Rand Journal of Economics*, 37, 887-909.
- [2] Basu, Amiya, Tridib Mazumdar and S.P. Raj (2003), "Indirect Network Externality Effects on Product Attributes," *Marketing Science*, 22, 209-221.
- [3] Bajari, Patrick, Lanier Benkard and Jon Levin (2007), "Estimating Dynamic Models of Imperfect Competition," *Econometrica*, 75 (5), 1331-1370.
- [4] Benkard, Lanier (2004), "A Dynamic Analysis of the Market for Wide-Bodied Commercial Aircraft," *Review of Economic Studies*, 71 (3), 581-611.
- [5] Blakely, Ryhs (2007), "Halo 3 Blows Away Rivals to Set Record," *Times Online*, Times Newspapers, October 4, 2007 .
- [6] Cabral, Luis (2007), "Dynamic Price Competition with Network Effects," *Working Paper*, New York University.
- [7] Carranza, Juan Esteban (2006), "Estimation of Demand for Differentiated Durable Good," manuscript, University of Wisconsin-Madison.
- [8] Chen, Jiawei, Ulrich Doraszelski and Joseph E. Harrington (2007), "Avoiding Market Dominance: Product Compatibility in Markets with Network Effects," Working Paper, UC Irvine.
- [9] Chou, Chien-fu, and Oz Shy (1990), "Network Effects Without Network Externalities," *International Journal of Industrial Organization*, 8, 259-270.
- [10] Church, Jeffrey, and Neil Gandal (1993), "Complementary Network Externalities and Technological Adoption," *International Journal of Industrial Organization*, 11, 239-260.
- [11] Clements, Matthew and Hiroshi O'Hashi (2005), "Indirect Network Effects and the Product cycle, Video Games in the US, 1994-2002," *Journal of Industrial Economics*, 53, 515-542.
- [12] Cobb, Jerry (2003), "Video-game makers shift strategy," 4:37 PM EST CNBC, <http://moneycentral.msn.com/content/CNBCTV/Articles/TVReports/P47656.asp>.
- [13] Coughlan, Peter J. (2001a), "Competitive Dynamics in Home Video Games (A): The Age of Atari", *Harvard Case*, HBS 9-701-091.
- [14] Coughlan, Peter J. (2001b), "Competitive Dynamics in Home Video Games (I): The Sony Playstation," *Harvard Case*, HBS 9-701-099.

- [15] Coughlan, Peter J. (2001c), “Competitive Dynamics in Home Video Games (K): Playstation vs. Nintendo64,” *Harvard Case*, HBS 9-701-101.
- [16] Coughlan, Peter J. (2001), “Competitive Dynamics in Home Video Games (I): The Sony Playstation,” *Harvard Case*, HBS 9-701-099.
- [17] Erdem, Tulin, Susumu Imai and Michael Keane (2003), “A Model of Consumer Brand and Quantity Choice Dynamics under Price Uncertainty,” *Quantitative Marketing and Economics*, 1, 5-64.
- [18] Evans, David S. (2003), “The Antitrust Economics of Multi-Sided Platform Markets,” *Yale Journal on Regulation*, 20, 325-381.
- [19] Evans, David S. and Richard Schmalensee (2007), “The Industrial Organization of Markets with Two-Sided Platforms,” *Competition Policy International*, 3, 151-179.
- [20] Farrell, Joseph, and Paul Klemperer (2006), “Coordination and Lock-In: Competition with Switching Costs and Network Effects,” *Handbook of Industrial Organization*, Vol. III (forthcoming).
- [21] Farrel J. and Saloner G. (1986), “Installed Base and Compatibility,” *American Economic Review*, Vol. 76, 940-955.
- [22] Gandal, Neil, 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*, 31-1.
- [23] Gallagher, Dan (2007) “Video game sales blow away forecasts,” *MarketWatch*, October 18, 2007 SAN FRANCISCO
- [24] Gordon, Brett R. (2006), “A Dynamic Model of Consumer Replacement Cycles in the PC Processor Industry,” Working Paper, Carnegie Mellon.
- [25] Gowrisankaran, Gautam and Marc Rysman (2006), “Dynamics of Consumer Demand for New Durable Goods,” Working Paper, Washington University, St. Louis.
- [26] Gupta, Sachin, Dipak C. Jain and Mohanbir S. Sawhney (1999), “Modeling the Evolution of Markets with Indirect Network Externalities: An Application to Digital Television,” *Marketing Science*, 18, 396-416.
- [27] Horsky, Dan (1990), “A Diffusion Model Incorporating Product Benefits, Price, Income, and Information,” *Marketing Science*, 9, 342-365.
- [28] Jenkins, Mark, Paul Liu, Rosa L. Matzkin and Daniel L. Mcfadden (2004), “The Browser War – Econometric Analysis of Markov Perfect Equilibrium in Markets with Network Effects,” Working Paper, Stanford University.

- [29] Karaca-Mandic, Pinar (2004), "Estimation and Evaluation of Externalities and Complementarities," Ph.D. dissertation, University of California, Berkeley.
- [30] Katz, Michael L., and Carl Shapiro (1985), "Systems Competition and Network Effects," *Journal of Economic Perspectives*, 8 (2), 93-115.
- [31] Katz, Michael L., and Carl Shapiro (1986a), "Technology Adoption in the Presence of Network Externalities," *The Journal of Political Economy*, 94, 822-841.
- [32] Katz, Michael L., and Carl Shapiro (1986b), "Product Market Compatibility Choice in a Market with Technological Progress" *Oxford Economic Papers, Special Issue on the New Industrial Economics*, 38, 146-165.
- [33] Katz, Michael L., and Carl Shapiro (1992), "Product Introduction with Network Externalities," *The Journal of Industrial Economics*, 94, 822-841.
- [34] Koski, Heli and Tobias Kretschmer (2004), "Survey on Competing in Network Industries: Firm Strategies, Market Outcomes, and Policy Implications," *Journal of Industry, Competition and Trade*, Bank Papers, 5-31.
- [35] Kunii, Irene (1998), "The Games Sony Plays; can it use Playstation to win mastery of cyberspace?" *Business Week*, New York: June 15, 1998, Iss. 3582, 34.
- [36] Liu, Hongju (2007), "Dynamics of Pricing in the Video Game Console Market: Skimming or Penetration?" Working Paper, University of Chicago.
- [37] Nair, Harikesh (2006), "Intertemporal Price Discrimination with Forward-looking Consumers: Application to the US Market for Console Video-Games," forthcoming *Quantitative Marketing and Economics*.
- [38] Nair, Harikesh, Pradeep Chintagunta and J.P. Dube (2004), "Empirical Analysis of Indirect Network Effects in the Market for Personal Digital Assistants," *Quantitative Marketing and Economics*, 2, 23-58.
- [39] Meese, Richard A. and Kenneth Rogoff (1982), "Empirical Exchange Rate Models of the Seventies, do they fit out of sample?" *Journal of International Economics*, 14, 3-24.
- [40] Pesendorfer, Martin and Phillip Schmidt-Dengler (2006), "Asymptotic Least Squares Estimators for Dynamic Games," Working Paper, London School of Economics
- [41] Petrin, Amil and Kenneth Train (2007), "Control Function Corrections for Omitted Attributes in Differentiated Product Models," Working Paper, University of Minnesota.
- [42] Rigdon, Joana (1996), "Nintendo 64 revitalizes slumping video-game market," *Wall Street Journal*, Dec. 17.

- [43] Rogoff, Kenneth (2007), "Comment on Exchange Rate Models Are Not as Bad as You Think by Charles Engel, Nelson C. Mark, and Kenneth D. West," For NBER Macroeconomics Annual 2007, Daron Acemoglu, Kenneth Rogoff and Michael Woodford, eds.
- [44] Ryan, Stephen and Catherine Tucker (2007), "Heterogeneity and the Dynamics of Technology Adoption," Working Paper, MIT.
- [45] Rust, John (1987), "Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher," *Econometrica*, 55 (5), 999-1033.
- [46] Rysman, Marc (2004), "Competition Between Networks: A Study of the Market for Yellow Pages," *Review of Economic Studies*, 71, 483-512.
- [47] Rysman, Marc (2004), "Empirics of Antitrust in Two-Sided Markets," *Competition Policy International*, 3, 483-512.
- [48] Shankar, Venkatesh and Barry Bayus (2003), "Network Effects and Competition: An Empirical Analysis of the Home Video Game Market," *Strategic Management Journal*, 24, 375-384.
- [49] Shapiro, Carl and Hal R. Varian (1999), "The Art of Standards Wars," *California Management Review*, 41, 8-32.
- [50] Williams, Dmitri (2007), "Structure and Competition in the US Home Video Game Industry," *International Journal on Media Management*, 4, 41-54.

Appendix A: Equilibrium Provision of Software

In this appendix, we illustrate how we can derive the hardware demand model based on tastes for variety of software. We use a CES model of preferences for software and assume a spot market of monopolistically competitive software suppliers.

After purchasing a hardware platform j , a consumer i purchases an assortment of compatible software each period, $x_{it} = (x_{i1t}, \dots, x_{in_{jt}})'$, by maximizing her software utility subject to the budget constraint:

$$\max_{\{x_1, \dots, x_{n_{jt}}\}} U_{ij}^{SW}(x_{i1}, \dots, x_{in_{jt}}, z_i) \quad s.t. \quad \sum_{k=1}^{n_{jt}} \rho_k x_{ik} + z_i = I_i - p_{jt} Q_{ijt}$$

where p_{jt} is the price of hardware standard j in period t , and Q_{ijt} indicates whether the consumer also purchases hardware that period (it is zero if they adopted hardware at some time period prior to t). The term z_i is a numeraire capturing expenditures on other goods, and ρ_k is the price of software k . We use CES preferences to model the consumer's utility for software:

$$U_{ij}^{SW}(x_{i1}, \dots, x_{in_{jt}}, z_i) = \left(\sum_{k=1}^{n_{jt}} x_{ik}^{\frac{1}{b}} \right)^{\frac{1}{a}} + \alpha z_i, \quad a \geq 1, b > 1.$$

The corresponding individual demand for software k is:

$$x_{kt}^* = (ab\alpha)^{\frac{ab}{1-ab}} \rho_k^{\frac{b}{1-b}} \left(\sum_{l=1}^{n_{jt}} \rho_l^{\frac{1}{1-b}} \right)^{\frac{ab-b}{1-ab}}.$$

Turning to the software supply side, we assume that consumers derive utility from software for only one period. Hence, a software firm earns profits on a software product for only one period. Let y_{jt+1} represent the installed base of consumers that have adopted hardware standard j prior to period $(t+1)$. This installed base represents the potential demand at time t for a manufacturer of software compatible with hardware standard j . Each software title is treated as a separate firm. The profit function for a software firm k active in period t producing a software title compatible with hardware j is:

$$\pi_{kt} = (\rho_k - c) y_{jt+1} x_{kt}^* - F$$

where F is the fixed development cost and c is the marginal cost. The marginal costs consist of both royalties to the manufacturer and physical production costs (e.g. CDs and cartridges for Sony and Nintendo respectively). Since software firms are assumed to be ex ante identical,

there exists a symmetric price equilibrium in which each firm sets prices as follows:

$$\rho = \alpha c.$$

This symmetry also allows us to simplify the demand function:

$$x_{kt}^* = (ab\alpha\rho)^{\frac{ab}{1-ab}} n_{jt_t}^{\frac{ab-b}{1-ab}}.$$

Under free entry, the equilibrium number of software firms, n_{jt} can be characterized by the installed base as follows:

$$\log(n_{jt}) = \kappa + \lambda \log(y_{jt+1}) \quad (9)$$

where $\kappa = \frac{ab-1}{ab-b} \frac{\log(\rho-c)}{F(ab\alpha\rho)^{\frac{ab}{ab-1}}} = \frac{ab-1}{ab-b} \frac{\log(c(\alpha-1))}{F(ab\alpha^2c)^{\frac{ab}{ab-1}}}$ and $\lambda = \frac{ab-1}{ab-b}$.

We now derive the aggregate sales of software for each standard. Total software sales will be important in determining the software royalties that accrue to each hardware firm. We can substitute (9) to express individual demand for software k as follows:

$$\begin{aligned} x_{kt}^* &= (ab\alpha\rho)^{\frac{ab}{1-ab}} n_{jt}^{\frac{ab-b}{1-ab}} \\ &= (ab\alpha^2c)^{\frac{ab}{1-ab}} \left(\exp(\kappa) y_{jt+1}^\lambda \right)^{\frac{-1}{\lambda}} \\ &= (ab\alpha^2c)^{\frac{ab}{1-ab}} \left(\exp\left(-\frac{\kappa}{\lambda}\right) \right) y_{jt+1}^{-1} \end{aligned}$$

We then obtain the corresponding Aggregate Demand for software k :

$$\begin{aligned} X_{kt}^* &= x_{kt}^* y_t \\ &= (ab\alpha^2c)^{\frac{ab}{1-ab}} \exp\left(-\frac{\kappa}{\lambda}\right) \end{aligned}$$

Finally, we obtain total software sales for the standard j :

$$\begin{aligned} Q_{jt} &= \sum_{k=1}^{n_{jt}} X_{kt}^* \\ &= N_{jt} \left[(ab\alpha^2c)^{\frac{ab}{1-ab}} \exp\left(-\frac{\kappa}{\lambda}\right) \right] \\ &= \left[\exp(\kappa) y_{jt+1}^\lambda \right] \left[(ab\alpha^2c)^{\frac{ab}{1-ab}} \exp\left(-\frac{\kappa}{\lambda}\right) \right] \\ &= \exp\left(\frac{\kappa(1-\lambda)}{\lambda}\right) (ab\alpha^2c)^{\frac{ab}{1-ab}} y_{jt+1}^\lambda. \end{aligned}$$

We can therefore estimate the elasticity of total software sales with respect to the installed

base as follows:

$$\log(Q_{jt}) = \phi + \lambda \log(y_{jt+1}) .$$

Appendix B: Forward-Simulation of the Consumers' Choice-Specific Value Functions

We outline the procedure for using the first-stage estimates of the consumers's choice strategy, (8), the console firms' pricing strategies, (7), and the software supply, (6), to forward-simulate the consumers' choice-specific value functions.

Conditional on the first-stage estimates and some initial state, y_0 , we can simulate histories of all variables affecting the consumers' payoffs. For any period t with beginning-of-period installed base y_t , we draw recursively as follows:

$$\begin{aligned}
 \xi_{jt} &\sim N(0, 1), & (j = 1, \dots, J) \\
 p_{jt}|y_t, \xi_t &= \exp(\mathcal{P}_j(y_t; \hat{\theta}_p) + \hat{\lambda}\xi_{jt}), \\
 \mu_{jt}|y_t, \xi_t &= \mathcal{L}_j(y_t, \xi_t; \hat{\theta}_\mu), \\
 s_{jt}|\mu_t &= \frac{\exp(\mu_{jt})}{1 + \sum_{k=1}^J \exp(\mu_{kt})}, \\
 y_{j,t+1}|y_t, s_t &= f_j(y_t, \xi_t) = y_{jt} + \left(1 - \sum_{k=1}^J y_{kt}\right) s_{jt}, \\
 n_{jt}|y_{j,t+1} &= \exp(\mathcal{H}_j(y_{j,t+1}; \hat{\theta}_n)).
 \end{aligned}$$

In this manner, we can draw a sequence of states, $\{y_t, \xi_t\}_{t=0}^T$, and corresponding prices, number of software titles, and market shares.

Choice-specific value functions We first compute the software value functions. We assume the current software utility is given by

$$u_j(y_{j,t+1}) = \gamma \exp(\mathcal{H}_j(y_{j,t+1}; \hat{\theta}_n)) = \gamma n_{jt}.$$

For any initial installed base y_0 , we draw a sequence of states $\{y_t^{(r)}, \xi_t^{(r)}\}_{t=0}^T$ and a sequence of corresponding software titles, $\{n_t^{(r)}\}_{t=0}^T$. Repeating this process R times, we calculate the simulated expected PDV of software at state $y = y_0^{(r)}$,

$$\mathcal{W}_j(y; \Lambda, \hat{\Theta}) = \frac{1}{R} \sum_{r=1}^R \left(\sum_{t=0}^T \beta^t \gamma n_{jt}^{(r)} \right).$$

The consumers' choice-specific value functions from adopting standard j can then be calculated as

$$\mathcal{V}_j(y, \xi, p; \Lambda, \hat{\Theta}) = \delta_j + \mathcal{W}_j(f(y, \xi); \Lambda, \hat{\Theta}) - \alpha p_j + \psi \xi_j.$$

Here, $\Lambda = (\delta, \alpha, \gamma, \psi)$ is a vector containing all the stage 2 preference parameters to be estimated. Note that T needs to be chosen large enough such that β^T is sufficiently small.

Value of waiting First, we define the expected per-period utility of a consumer who has not adopted at the beginning of period t , conditional on y_t , p_t , and ξ_t :

$$\mathcal{U}(y_t, \xi_t) = s_{0t}\mathbb{E}(\epsilon_{0t}|0) + \sum_{j=1}^J s_{jt}(\delta_j + \gamma n_{jt} - \alpha p_{jt} + \xi_{jt} + \mathbb{E}(\epsilon_{jt}|j)).$$

In this equation, s_t , p_t , and n_t are the choice probabilities, prices, and number of software titles as implied by the first-stage estimates, conditional on the current states y_t and ξ_t . Furthermore, $\mathbb{E}(\epsilon_{jt}|j) = -\log(s_{jt})$ is the expected value of the Type I Extreme Value random utility component, given that choice j is optimal.

Next, we define m_{0t} as the probability that a consumer has not adopted one of the hardware standards prior to period t . Note that $m_{01} = 1$, because we want to calculate the value of waiting in period $t = 0$. Thereafter ($t > 1$), m_{0t} evolves according to

$$m_{0t} = s_{0,t-1}m_{0,t-1}.$$

m_{jt} denotes the probability that a consumer has adopted standard j prior to period t . $m_{j1} = 0$, and for $t > 1$,

$$m_{jt} = m_{j,t-1} + s_{j,t-1}m_{0,t-1}.$$

We now draw some sequence of states, $\{y_t^{(r)}, \xi_t^{(r)}\}_{t=0}^T$, with initial conditions $(y, \xi) = (y_0^{(r)}, \xi_0^{(r)})$. Given a corresponding sequence of $m_{0t}^{(r)}$ and $m_{jt}^{(r)}$, define

$$\mathcal{J}^{(r)} = \sum_{t=1}^T \beta^t \left(m_{0t}^{(r)} \mathcal{U}(y_t^{(r)}, \xi_t^{(r)}) + \sum_{j=1}^J m_{jt}^{(r)} (\gamma n_{jt}^{(r)}) \right).$$

$\mathcal{J}^{(r)}$ is the expected present discounted value from waiting, given that the market evolves according to $\{y_t^{(r)}, \xi_t^{(r)}\}_{t=0}^T$. Averaging over R draws, we obtain the expected value from waiting, conditional on $(y, \xi) = (y_0^{(r)}, \xi_0^{(r)})$:

$$\mathcal{V}_0(y, \xi; \Lambda, \hat{\Theta}) = \frac{1}{R} \sum_{r=1}^R \mathcal{J}^{(r)}.$$

Table 1: Descriptive Statistics

	Console	Mean	SD	Min	Max
Sales	Playstation	275,409	288,675	26,938	1,608,967
	Nintendo	192,488	201,669	1,795	1,005,166
Price	Playstation	119.9	30.3	55.7	200.6
	Nintendo	117.6	33.9	50.3	199.9
Game Titles	Playstation	594.2	381.1	3	1,095
	Nintendo	151.2	109.9	1	281

Table 2: Model Fit for Different Specifications

Model	Log-Likelihood	BIC
1) Linear, ξ , 1-comp	-187.88	679.40
2) Linear, time ($t < 60$), 1-comp	-150.11	620.97
3) Quadratic, time ($t < 60$), 1-comp	-79.38	514.01
4) Quadratic, time ($t < 60$), 2-comp	-79.38	522.28
5) Quadratic, time ($t < 60$), 1-comp, PPIs in prices	-63.54	507.69
6) Quadratic, time ($t < 60$), 1-comp, exchange rate in prices	-25.43	422.94
7) Quadratic, time ($t < 60$), 1-comp, exchange rate and PPI's in prices	-7.53	412.79

Table 3: First Stage Estimates: Pricing Policies \mathcal{P}_j

	Sony		Nintendo	
	Estimate	SE	Estimate	SE
Intercept	2.653	1.523	-1.592	4.051
y_{Sony}	-3.266	1.785	-0.861	2.490
y_{N64}	1.889	1.499	0.666	0.813
y_{Sony}^2	-0.320	0.218	-0.142	0.282
y_{N64}^2	0.186	0.130	0.077	0.070
Time (< 64)	0.005	0.003	0.000	0.002
Jan	0.102	0.040	-0.054	0.016
Feb	0.132	0.119	0.005	0.104
Mar	0.118	0.053	-0.082	0.214
Apr	0.136	0.057	-0.046	0.046
May	0.045	0.032	-0.055	0.037
Jun	-0.012	0.003	-0.011	0.028
Jul	0.000	0.026	-0.052	0.037
Aug	-0.001	0.012	-0.032	0.021
Sep	-0.049	0.074	-0.088	0.021
Oct	0.011	0.048	-0.015	0.039
Nov	-0.026	0.019	-0.002	0.025
PPI 1	-0.432	0.070	-0.285	0.154
PPI 2	-0.546	0.351	-0.608	0.816
PPI 3	-0.471	0.769	0.227	0.807
Exchange rate (3 month lag)	0.434	0.509	0.397	0.159
Exchange rate (7 month lag)	-4.493	1.118	1.390	8.096
$\log(\lambda)$	-4.511	0.123	-4.511	0.123

Table 4: First Stage Estimates: Log-odds of Market Shares, \mathcal{L}_j

	Sony		Nintendo	
	Estimate	SE	Estimate	SE
Intercept	-13.826	2.539	-1.141	0.744
y_{Sony}	1.456	4.669	0.065	0.012
y_{N64}	-7.670	5.884	-1.740	0.196
y_{Sony}^2	-0.334	0.432	-1.304	0.305
y_{N64}^2	-0.684	0.571	-1.820	0.335
Time (< 64)	-0.003	0.008	-2.176	0.469
Jan	-1.367	0.154	-1.950	0.528
Feb	-1.345	0.159	-1.625	0.495
Mar	-1.813	0.475	-1.682	0.525
Apr	-2.442	0.351	-1.674	0.397
May	-2.596	0.463	-1.268	0.402
Jun	-1.947	0.395	-1.713	0.215
Jul	-1.805	0.470	-0.789	0.335
Aug	-1.871	0.392	-0.288	0.070
Sep	-1.496	0.426	0.085	0.135
Oct	-1.644	0.199	-0.406	0.150
Nov	-0.781	0.120	0.084	0.173
ξ_{Sony}	-27.656	5.216	0.103	0.061
ξ_{N64}	0.844	4.936	-0.220	0.083
ξ_{Sony}^2	-13.383	7.554	0.547	0.107
ξ_{N64}^2	-0.660	0.416	-0.545	0.119

Table 5: First Stage Estimates: Equilibrium Game Provision, \mathcal{H}_j

	Estimate	SE
κ_{Sony}	-16.220	2.042
$\kappa_{Nintendo}$	-24.349	1.992
λ_{Sony}	1.369	0.126
$\lambda_{Nintendo}$	1.810	0.126

Table 6: Second Stage Parameter Estimates

	Model 3		Model 7	
	Estimate	SE	Estimate	SE
δ_{Sony}	-1.21	0.89	-1.119	0.971
δ_{N64}	-1.34	0.87	-1.119	1.093
α	-1.94	0.52	-1.923	0.46
Time (< 60)	-0.04	0.01	-0.049	0.028
γ ($n_{jt}/1000$)	0.09	0.04	0.09	0.04
ψ (s.d. of ξ_{jt})	0.05	0.09	0.028	1.95

Notes: Model 7 uses PPI's and exchange rates as IV's in first stage.
 $\beta = 0.9$, no. simulations = 60

Table 7: Predicted One-Firm Concentration Ratios

<i>Model Predictions:</i>				
<i>Symmetric Case (Parameter Estimates for Sony)</i>				
Scale Factor for γ	0.25	0.50	0.75	1.00
C_1	0.502	0.503	0.506	0.833
Discount Factor (β)	0.6	0.7	0.8	0.9
C_1^a	0.502	0.502	0.502	0.833
C_1^b	0.501	0.501	0.503	0.833
<i>Model Predictions:</i>				
<i>Estimated Parameter Values</i>				
Scale Factor for γ	0.25	0.50	0.75	1.00
C_1	0.594	0.581	0.552	0.827
Discount Factor (β)	0.6	0.7	0.8	0.9
C_1^a	0.597	0.594	0.587	0.827
C_1^b	0.566	0.564	0.553	0.827

Note: ^(a) All estimated model parameters were obtained for $\beta = 0.9$. ^(b) Predictions where the model parameters were re-estimated for each consumer discount factor, β .

Table 8: Predicted Degree of Tipping at Estimated Parameter Values ($\beta = 0.9$)

Scale Factor for γ	0.25	0.50	0.75
ΔC_1	0.233	0.246	0.275
Discount Factor (β)	0.6	0.7	0.8
ΔC_1^a	0.230	0.233	0.240
ΔC_1^b	0.261	0.263	0.274

Note: The table displays the increase in market concentration relative to a specific counter-factual model, where either the marginal utility of software, γ , is scaled, or a different consumer discount factor β is chosen. ^(a) All estimated model parameters were obtained for $\beta = 0.9$. ^(b) Predictions where the model parameters were re-estimated for each consumer discount factor, β .




Figure 1: In-Sample Fit: Prices




Figure 2: In-Sample Fit: Log-Odds Ratios




Figure 3: In-Sample Fit: Provision of Games




Figure 4: Symmetric competition: Equilibrium pricing policies and price paths. Consumer's software utility coefficient is scaled by different factors. The expected price paths are shown conditional on $y_{T1} \geq y_{T2}$ at the end of period $T = 25$. Marginal production costs are indicated by horizontal lines.




Figure 5: Symmetric competition: Expected state evolution and distribution of shares in the installed base after 25 months. Consumer's software coefficient is scaled by different factors.

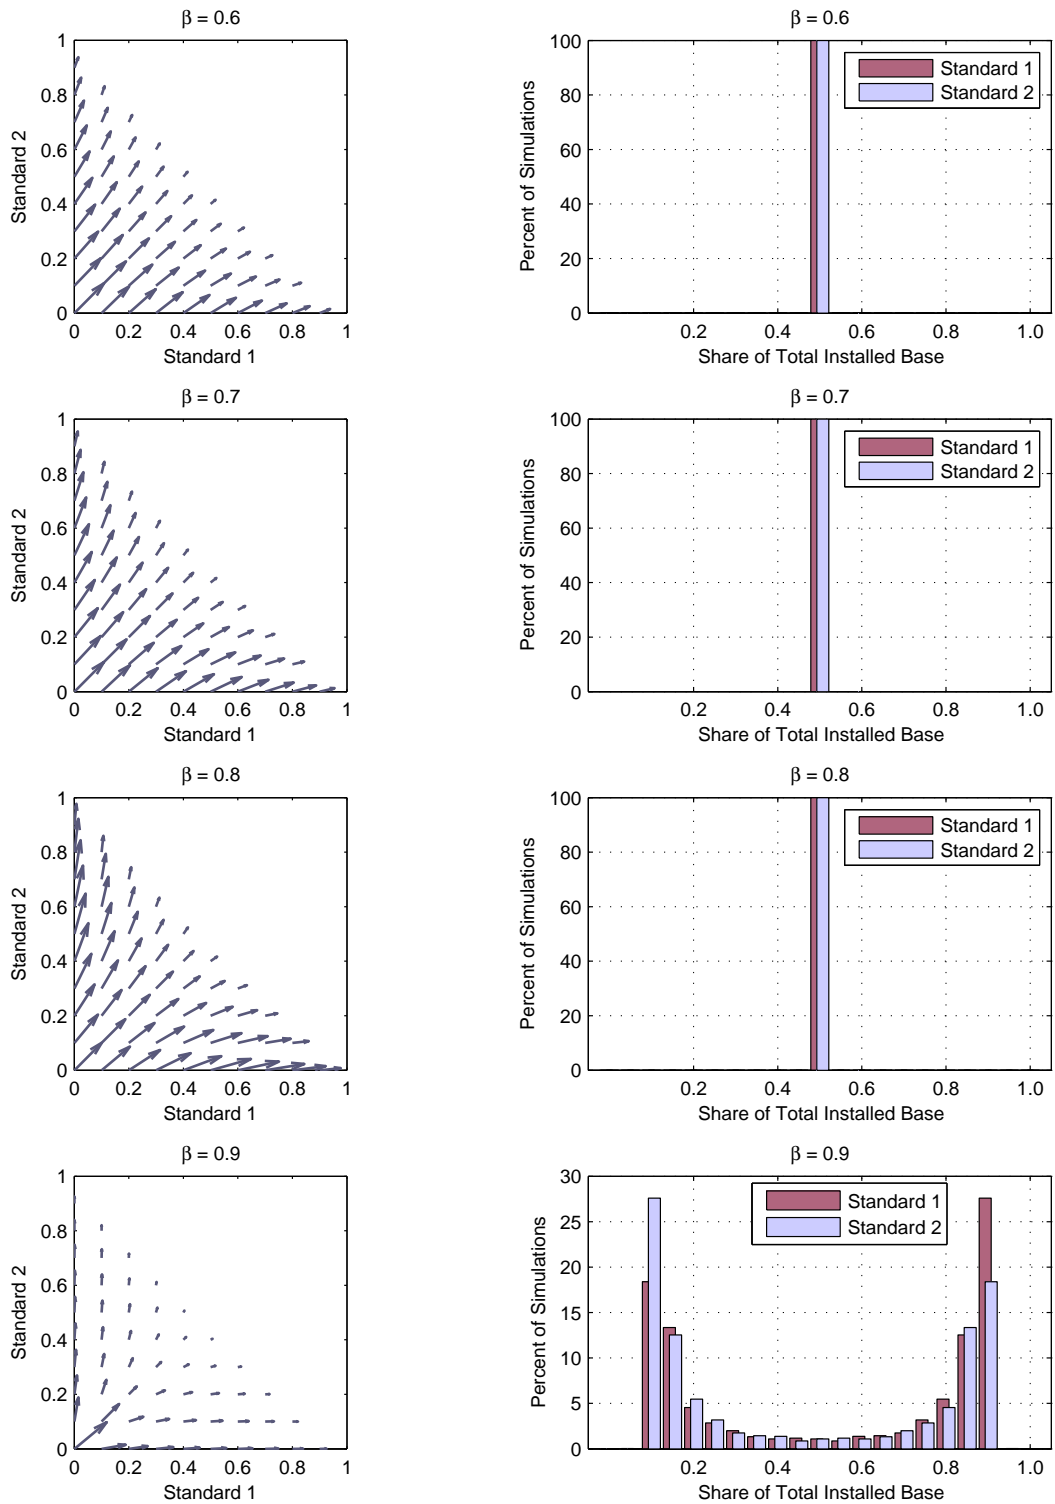


Figure 6: Symmetric competition: Equilibrium pricing policies and price paths for different consumer discount factors (β).




Figure 7: Predictions from estimated parameter values: Expected state evolution and distribution of shares in the installed base after 25 months. Consumer's software coefficient is scaled by different factors.




Figure 8: Predictions from estimated parameter values: Equilibrium pricing policies and price paths for different consumer discount factors (β).