Estimating the Innovator's Dilemma: Structural Analysis of Creative Destruction

Mitsuru Igami*

May 30, 2013

Abstract

Theories predict cannibalization between existing and new products delays incumbents' innovation, whereas preemptive motives accelerate it, and incumbents' cost (dis)advantage relative to entrants would further reinforce these tendencies. To empirically assess these three forces, I develop and estimate a dynamic oligopoly model using a unique panel dataset of hard disk drive (HDD) manufacturers (1981–98). The results suggest that despite strong preemptive motives and a substantial cost *advantage* over entrants, incumbents are reluctant to innovate because of cannibalization, which can explain at least 51% of the incumbent-entrant timing gap. I then assess hypothetical policy interventions including broad patents and trade barriers, and find the industry's welfare trajectory difficult to outperform, as Schumpeter (1942) conjectured.

Keywords: Creative Destruction, Dynamic Oligopoly, Structural Estimation

*Yale Department of Economics. E-mail: mitsuru.igami@yale.edu. I thank my dissertation advisers at UCLA: Daniel Ackerberg, Hugo Hopenhayn, Edward Leamer, Mariko Sakakibara, Connan Snider, and Raphael Thomadsen. For suggestions, I thank Ron Adner, Lanier Benkard, Steven Berry, Michael Dickstein, Pinelopi Goldberg, Jinyong Hahn, Philip Haile, Phillip Leslie, Rosa Matzkin, Matthew Mitchell, Ichiro Obara, Ariel Pakes, Marc Rysman, Kosuke Uetake, Yong Hyeon Yang, and Mark Zbaracki, as well as seminar participants at IIOC, TADC, REER, AEA, and various universities. I thank Clayton Christensen for encouraging a new approach to the innovator's dilemma, Minha Hwang for sharing engineering expertise, and James Porter, late editor of *DISK/TREND Reports*, for sharing industry knowledge and the reports. Financial support from the Nozawa Fellowship, the UCLA CIBER, and the Dissertation Year Fellowship is gratefully acknowledged. An earlier version of the paper received the Best Student Paper Award at *the 11th Annual REER* at Georgia Tech and the Dreze Award for the best paper by a PhD student at UCLA Anderson.

1 Introduction

"In the long run we are all dead,"[1] and firms and technologies are no exception. Netflix's movie download service has grown fast, whereas Blockbuster, a once-dominant DVD rental chain, filed for bankruptcy protection in 2010 after a reluctant pursuit of an online distribution service. Amazon is now selling everything from electronic books to disposable diapers, whereas Borders, America's number-two book retailer, liquidated its shops in 2011 after belated efforts to introduce its own electronic reader. These examples may seem extreme, but even when introducing a new technology is not too difficult, the old winners tend not to adapt well while the new entrants quickly become successful, according to a former CEO of Intel, the world's biggest chip maker (Grove 1996). Some incumbents never introduce a new technology/product despite shrinking demand for their existing products, a puzzling phenomenon called the innovator's dilemma (Christensen 1997). The timing of innovation in general and the incumbent-entrant timing gap in particular are important for both businesses and policymakers. Who innovates and survives better (and why) is a central question for individual firms. Depending on how incumbents' and entrants' incentives differ, competition and innovation policies will have different consequences. Understanding the determinants of the timing gap is the first step toward designing a pro-innovation competition policy (Bresnahan 2003). For these purposes, this paper empirically tests three theoretical determinants of incumbents' innovation.[2]

Why do incumbents delay innovation? Viewed from a microeconomic perspective, the determinants of innovation timing include (1) cannibalization, (2) different costs, (3) preemption, and (4) institutional environment (Hall 2004, Stoneman and Battisti 2010). First, because of cannibalization, the benefits of introducing a new product are smaller for incumbents than for entrants, to the extent that the old and new goods substitute for each other. By introducing new goods, incumbents are merely replacing their old source of profits, so Arrow (1962) calls this mechanism the "replacement effect." Second, incumbents may face higher costs of innovation because of organizational inertia. Economic theory, as well as case studies, suggest that as firms grow larger and older, their R&D efficiency diminishes (e.g., Schumpeter 1934);[3] although, *a priori*, hypothesizing that incumbency confers some advan-

[1] John Maynard Keynes, *A Tract on Monetary Reform* (1923), Ch. 3.

[2] Following Schumpeter's (1934) view, this paper uses the words "innovation" and "technology adoption" interchangeably, and avoids equating innovation with "invention," which may or may not constitute innovation in the economic sense. In the HDD industry, the technological roadmap and the future product standards (i.e., new concepts) are widely shared, whereas the actual commercialization and product development are left to individual firms' efforts. This paper takes the former (i.e., shared conceptualization) as exogenously given, and analyzes the latter (i.e., commercialization) as endogenous decisions.

[3] The existing literature suggests various reasons for incumbents' inertia, such as bureaucratization (Schumpeter 1934), hierarchy (Sah and Stiglitz 1986), loss of managerial control (Scherer and Ross 1990),

tages due to accumulated R&D capital is equally plausible (e.g., Schumpeter 1942). Hence, whether incumbents have a cost advantage or disadvantage is an open empirical question. Third, market structure dynamics play an important, countervailing role, as theories predict incumbents should innovate more aggressively than entrants to preempt potential rivals (e.g., Gilbert and Newbury 1982) under various oligopolistic settings. Finally, the impact of these three determinants will change under different institutional contexts, such as the rules governing patents and international trade. In total, these three competing forces (plus institutional contexts) determine innovation timing. Cannibalization delays incumbents' innovation, whereas preemptive motives accelerate it, and incumbents' cost (dis)advantage would further reinforce these tendencies. Given this tug of war between the three theoretical forces, I propose to explicitly incorporate them into a unified model and estimate the empirical importance of each.

Figure 1: Shifting Generations of Technology

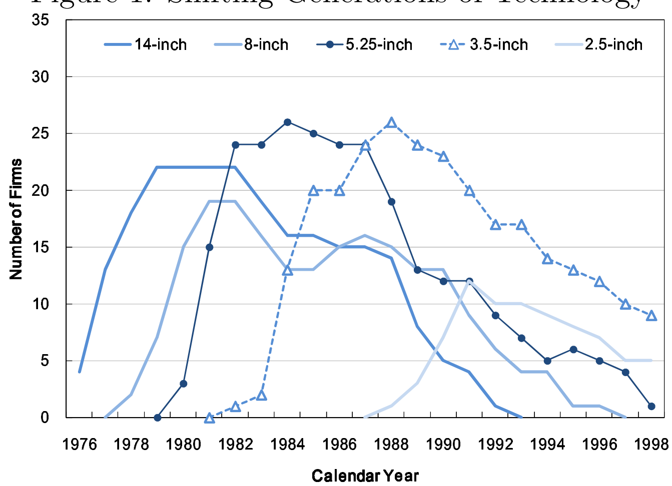

Note: Shipment-based recognition of firms. Major firms only.

The goal of this paper is to provide the first structural econometric account of creative destruction (Schumpeter 1942) and empirically quantify these competing forces behind the innovator's dilemma in the HDD industry, which is a highly relevant setting. As epitomized by Christensen (1997), this industry is a canonical case of creative destruction, where cohorts of firms come and go with the generational transitions of technologies (Figure 1). I construct a unique dataset from the industry publications, *DISK/TREND Reports* (1977–99), which record a comprehensive set of firms (both incumbents and potential entrants) for more than two decades. An HDD is a component of a personal computer (PC) that stores information. Desktop PCs used 5.25-inch HDDs during the 1980s, but 3.5-inch HDDs became popular

and informational, cognitive, or relationship reasons (Grove 1996, Christensen 1997).

during the 1990s, so those firms that exclusively manufactured 5.25-inch HDDs disappeared by the turn of the century.

The task of quantifying the three forces calls for a structural approach, because incentives to innovate are sensitive to an industry's technological and institutional context. I build and estimate a dynamic oligopoly model that explicitly incorporates cannibalization, heterogeneous innovation costs, and preemptive motives (dynamic strategic interactions), all of which endogenously determine the timing of innovation and the evolution of market structure. Then I measure the effects of the three forces by contrasting the outcome of the estimated model with those of three counterfactual simulations in which firms ignore each of these incentives. Finally, to study broader implications of the phenomenon, I simulate evolutions of the industry under two alternative institutional settings: (1) a broad patent regime and (2) a ban on foreign goods.[4]

My modeling and estimation approach proceeds in three steps. First, I estimate demand using a standard discrete choice model for differentiated goods (the old- and new-generation HDDs). That is, I let the data tell the degree of substitution between the old and new goods and hence of cannibalization. Second, I recover marginal costs of production, implied by the first-order conditions of static competition (multi-product Cournot). From these demand and cost estimates, I calculate the static equilibrium profits in every state of the industry, that is, given any number of active firms in the market. These profit estimates embody the relationship between market structure and profitability and hence give a preliminary indication of the extent to which preemption motivates firms' introduction of a new technology/product. Third, I feed these static (period) profits into a dynamic model to estimate the sunk costs of innovation. The model features two types of firms, incumbents and entrants, so I estimate the sunk costs separately for each type. I explicitly incorporate firms' dynamic discrete choice between entering, exiting, continuing operation with the old product, or introducing the new product. I fully incorporate preemptive motives, because firms interact strategically and are forward-looking with rational expectations over the endogenous evolution of market structure. To reflect the uniqueness of the technological transition in HDDs as well as the computer industry's ever-changing nature, I make my model non-stationary, allowing demand, costs, and hence value and policy functions to change over time.

Conceptually, this third step is simple. I invoke a revealed preference argument and employ maximum likelihood estimation (MLE), with nested fixed-point (NFXP) algorithm, to find the sunk cost parameter values that would best rationalize the actual innovation and entry/exit behaviors in the data.[5] Practically, however, this procedure poses two technical

[4]Appendix reports additional simulation results for theoretical benchmarks and two more policy experiments concerning competition and innovation (entry deregulation and R&D subsidies for incumbents).

[5]I have chosen to use NFXP because this paper studies non-stationary industry dynamics in a single, global

challenges. One is the possibility of multiple equilibria. I address this issue through parsimonious modeling (a small state space and choice sets, and period-by-period, state-by-state solutions) to achieve uniqueness in some industry states, while relying on equilibrium selection in other cases (see section 3.1 for details). The other problem is the computational burden of calculating the equilibrium strategies and values, for each of the several thousand industry states, in each of the 18 years, for each set of candidate parameter values. I address this issue by coding the most computationally expensive routines (the calculation of expected values) in the C language within the MATLAB platform.

The estimation and simulation results suggest that despite strong preemptive motives and a substantial cost *advantage* over entrants, incumbents are reluctant to innovate early because of cannibalization, which can explain at least 51% of the timing gap. Incumbents may lag behind entrants, despite their advantage in innovation costs, which suggests a substantial part of what researchers have previously understood as organizational inertia could potentially be reinterpreted as an effect of cannibalization. The results from the policy simulations highlight the pro-innovation effect of competition. For example, I find a ban on international trade discourages innovation and hurts consumers. However, social welfare sometimes improves under *anti*-competitive policies, such as broad patents. Ironically, welfare improves not through promotion of innovation, but through cost savings from preventing "excess" entry/innovation.

I have organized the rest of the paper as follows. The remainder of this section explains how this research contributes to the literature on innovation and industry dynamics. Section 2 summarizes the technological and institutional background of the HDD industry. Section 3 describes the model. Sections 4 and 5 explain the estimation procedure and results. Section 6 quantifies the three economic forces behind "the innovator's dilemma." Section 7 evaluates welfare consequences of different policies in innovation and competition. Section 8 concludes.

1.1 Related Literature

This paper uses a structural approach to study the process of creative destruction. As such, it builds on three bodies of literature: innovation, industry dynamics, and the estimation of dynamic games.

First, many papers have studied the relationship between competition and innovation, with mixed predictions and inconclusive evidence (see Gilbert [2006] and Cohen [2010] for detailed surveys). Arrow (1962) predicted an incumbent monopolist has less incentive to

market, which makes the use of other approaches difficult, such as two-step estimation and mathematical programming with equilibrium constraints (MPEC).

innovate than perfect competitors, because of the "replacement effect" (i.e., the substitution between the old and new technologies), against which others theorized the preemptive motive for an incumbent monopolist to innovate more aggressively than an entrant (e.g., Gilbert and Newbery 1982, Reinganum 1983, Fudenberg and Tirole 1986). Empirical findings are mixed. Gilbert (2006) attributed the inconclusiveness of evidence to (1) the failure to control for contingencies highlighted by theorists, (2) the presence of fundamental heterogeneities across industries, and (3) the shortcomings of data and methods. Cohen (2010) raised a fundamental critique that market structure was a function of innovation itself, and entry, competition, and innovation were simultaneously determined by more structural factors such as demand and technological opportunities. This paper addresses these issues by focusing on a specific high-tech industry, explicitly incorporating the technological context of the industry into a dynamic model.

Second, the co-evolution of technology and competition has played a central role in the studies of industry evolution, or market structure dynamics. Theoretical models and qualitative case studies constitute the bulk of this literature.[6] Recent growth models also incorporate innovation by both incumbents and entrants (e.g., Klette and Kortum 2004, Acemoglu and Cao 2010). However, quantitative empirical work has been rare, presumably because the drastic nature of the phenomenon poses challenges to both data collection and empirical methods. Given the context of this literature, I propose a formal empirical analysis by developing a structural model of the Ericson-Pakes class and applying it to data from the HDD industry, the canonical case of "disruptive innovation."[7]

Third, in terms of the empirical approach, the closest papers are Benkard (2004), Schmidt-Dengler (2006), and Goettler and Gordon (2011), which use full-solution methods in studying dynamic games of innovation. I extend the analytical scope to include entry/exit and investigate incumbent-entrant heterogeneity as well as market structure dynamics.

In short, this paper presents the first structural analysis of creative destruction, bridging the frameworks to analyze innovation and entry/exit in a simple model. This model is empirically tractable and motivated by the institutional background of the HDD industry, which the next section summarizes.

[6] Prominent models include those of Nelson and Winter (1978, 1982), Jovanovic (1982), Hopenhayn (1992), Ericson and Pakes (1995), Klepper (1996), and Sutton (1998). Numerous case studies record instances of the generational transitions of firms and technologies: Tushman and Anderson (1986), Henderson and Clark (1990), Henderson (1993). More recent papers formally model these generational transitions, such as Adner and Zemsky (2005) and Klepper and Thompson (2006).

[7] Christensen (1993, 1997), Chesbrough (1999), and King and Tucci (2002), among other management researchers, have studied this industry.

2 Industry and Data

This section describes the key features of the HDD industry and explains why it is particularly suitable for the study of innovation and industry evolution.

2.1 HDD: Canonical Case of Creative Destruction

The HDD industry provides a particularly fruitful example for the study of technological change and industry dynamics.

First, the HDD industry is the canonical example of "disruptive innovation." Multiple generations of technologies were born, matured, and died within a decade or two. A generation was defined by the diameter of disks used: 14-, 8-, 5.25-, 3.5-, and 2.5-inch (see Figure 1). The introduction of a new HDD of smaller diameters required a significant technological investment because each firm had to go through a process of trial and error in determining the adequate configuration of components, then build new assembly lines, and finally establish a reliable process for volume production. Along with each generation, a cohort of firms came and went, many of which delayed the adoption of a newer technology. Pooling the observations across five generations (4 transitions), Figure 2 plots the timing of innovation separately for incumbents (i.e., firms already active in the previous generation) and entrants (i.e., firms that appeared for the first time as the producer of new-generation HDDs). Only about a half of all incumbents ever innovated into a new generation. Even among those that did, their timing was approximately two years later than entrants. Those that never innovated gradually disappeared along with the shrinking demand for the old products. Changes in technology and market structure are pervasive in many industries, but the HDD market has witnessed one of the fastest, most unrelenting, and most easily measurable turnovers of products and firms.

Second, a detailed industry data book series, the *DISK/TREND Reports* (1977–99), is available for this industry. From the original reports, I construct a comprehensive panel of the world's HDD manufacturers by digitizing each firm-year observation. The sample period is long enough to capture five generations of technologies, two of which I analyze in detail.

Third, a high-tech manufacturing sector with rapid growth and innovation is precisely the type of industry that is most relevant to the discussion of pro-innovation public policies. Moreover, the HDD market's fairly competitive structure (a total of 178 unique firms over 23 years) and geographical outreach (firms from the Americas, Asia, and Europe compete in a global market) underline the potential generalizability of the findings.

My analysis of the HDD industry takes the developments in the global PC market (i.e.,

Figure 2: CDF of Adoption Timing

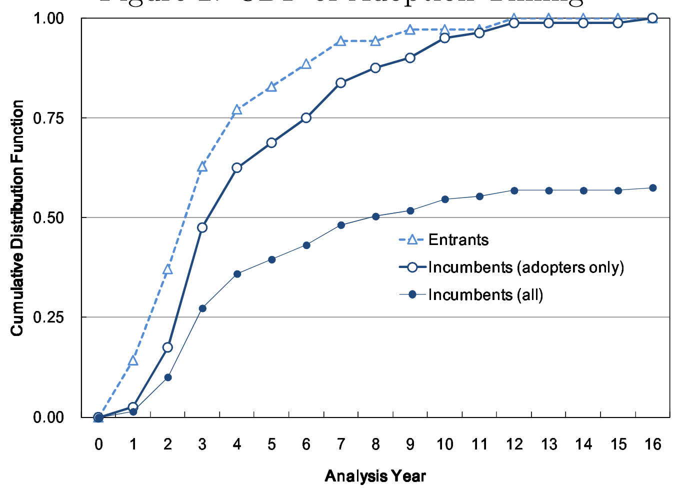

Note: Shipment-based recognition of technology adoption. Major firms only. Total of all diameters (14-, 8-, 5.25-, 3.5-, and 2.5-inch).

HDD's "downstream" industry) as given for the following reasons. First, the growth in PC demand was primarily driven by hobbyists during the 1980s and then by office automation and growing popularity among households during the 1990s. Second, the price and performance of the central processing unit (CPU) and operating system (OS) determined most of the cost and perceived quality of PCs, and hence the overall demand for PCs and their replacement purchase cycles. Although the quality improvement of HDDs contributed to the enhanced performance of PCs in terms of storage capacity, Intel and Microsoft ("Wintel") were perceived to be the leaders of the PC industry. Third, the market structure of PC makers is rather competitive, with more than 100 firms across the globe. As Table 1 shows, even the combined market share of the top five makers was less than 50%. Moreover, the vendors (brands) and manufacturers of PCs were often different; that is, many less well-known manufacturers made products for famous brands such as Compaq and Hewlett-Packard. Hence the market structure of actual manufacturers is less concentrated than what vendors' market share suggests.

Likewise, I do not explicitly model the developments in the "upstream" industry: HDD components, such as read-write heads, platters, and motors. Some HDD manufacturers make these components in-house, whereas others procure them from electronics parts makers. According to James Porter, the late editor of *DISK/TREND Reports*, there is no clear advantage/disadvantage concerning whether to make or to buy.

Table 1: Global PC Market Share by Units (%)

Rank	1981		1986		1991		1996	
1	Sinclair	23.9	IBM	12.3	IBM	11.4	Compaq	10.0
2	Apple	13.6	Commodore	11.4	Apple	9.1	IBM	8.6
3	Commodore	13.0	Apple	7.8	Commodore	8.3	Packard-Bell NEC	6.0
4	Tandy	12.8	Amstrad	5.9	NEC	5.8	Apple	5.9
5	Atari	4.5	NEC	5.3	Compaq	4.0	HP	N/A
Others		32.3		57.2		61.5		69.5

Note: Market share based on worldwide unit shipments.

Source: Gartner Dataquest, Wikipedia.

2.2 Data Source

I manually construct the panel data of 1,378 firm-year observations from *DISK/TREND Reports* (1977–99), an annual publication series edited by the HDD experts in Silicon Valley. I digitize each firm-year observation, which is accompanied by half a page of qualitative descriptions (on the characteristics of the firm, managers, funding, products, production locations, as well as major actions taken in that year, with their reasons) in the original publication. Not all information is amenable to quantitative analysis, but some of the firms' characteristics are suitable for regressions. For example, firms' age and size (in terms of revenues and profits, either company-wide or specifically for the HDD business) are readily codifiable.[8] Firms' organizational forms, regions of origin, and the initial HDD generations in which they started manufacturing, are also digitized as categorical variables.

An auxiliary dataset, also from *DISK/TREND Reports*, containing the prices and shipment quantities of HDDs, accompanies these panel data of firms. For each year, the reports record the average transaction price and total quantity for each of the generation-quality categories (5 generations and 14 quality levels in total).

Researchers in both economics and management repeatedly confirm the accuracy, relevance, and comprehensiveness of the record.[9]

2.3 Focus: Transition from 5.25- to 3.5-inch Generations

I analyze the technological transition from the 5.25- to 3.5-inch generations, which I will call the "old" and "new" generations henceforth. This subsample of the dataset spans 18 years

[8]However, age is not necessarily comparable across firms that had roots in different industries (e.g., manufacturers of card punchers, typewriters, automobile components, or coin laundries). Not all firms disclosed division-level revenue/profit information. For these reasons, I omit these variables in the following regressions.

[9]Christensen (1993, 1997); Lerner (1997); McKendrick, Donner, and Haggard (2000); and Franco and Filson (2006).

(1981–98) and 259 firm-years. I concentrate on these generations because they competed directly with each other in the desktop PC market. Although transitions between the other generations showed similar developments, 14-, 8-, and 2.5-inch HDDs were used in different segments of the computer industry, that is, 14-inch for mainframe computers, 8-inch for minicomputers, and 2.5-inch for notebook PCs. By focusing on 5.25- and 3.5-inch generations, I avoid confounding factors that might originate from diverging trends in different segments downstream. These two generations were also historically the most important of all generations in terms of volume and revenue.

Figure 3: Quantity, Price, Quality, and Industry Composition

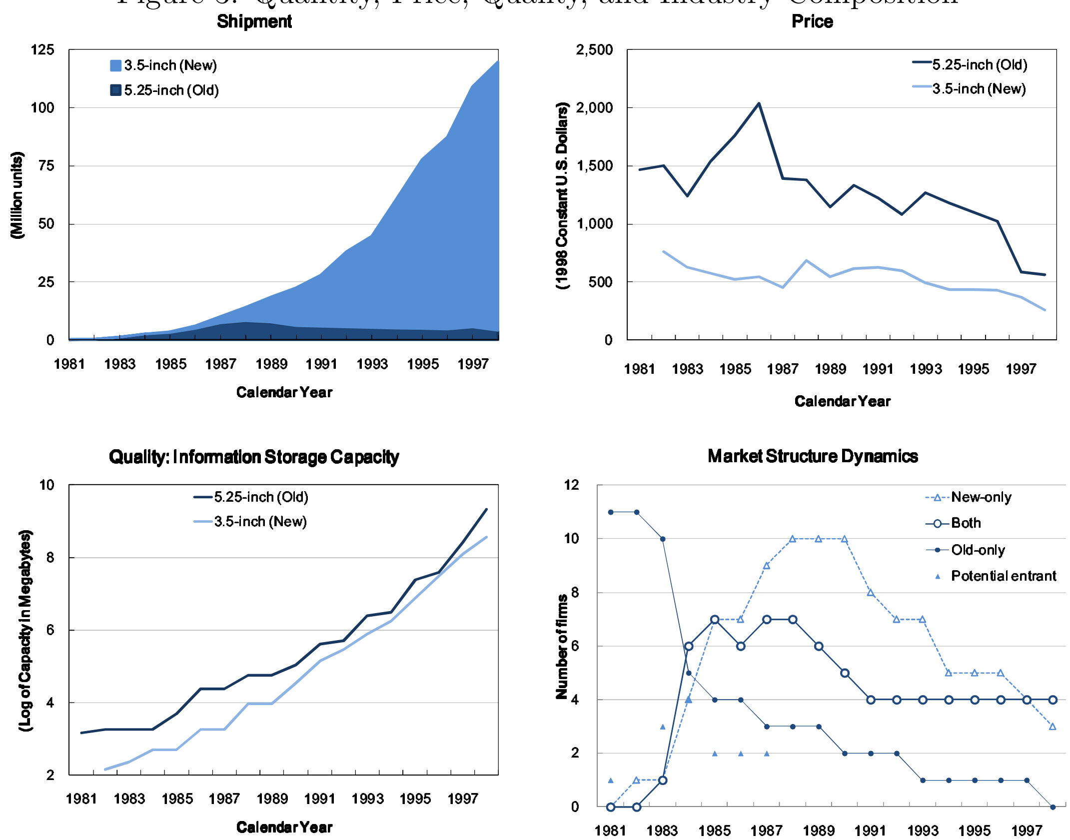

Note: Both 5.25- and 3.5-inch HDDs serve the same market, namely, desktop personal computers. Quality is measured by average capacity per unit for each generation. "Old-only" and "new-only" firms produce 5.25- and 3.5-inch HDDs, respectively. "Both" represents incumbents that adopted the new technology, hence producing both of the two generations. "Potential entrant" is identified by the announcement of product specifications (without actual shipment).

The two generations of HDD experienced a fast growth in volume and a steady decline in price (Figure 3, top). Over the years, the average quality (information storage capacity) of HDDs improved at an approximately constant rate (left bottom panel). These developments were typical of those in many computer-related industries. The right bottom panel shows the numbers of firms in four different states: (1) "old-only," (2) "both," (3) "new-only," and (4) "potential entrant." Incumbents start as (1) and become (2) upon the adoption of new technology. Entrants start as (4) and become (3) upon adoption (entry).

3 Modeling

This section presents my industry equilibrium model. The first subsection outlines the dynamic game of technology adoption and entry/exit. The second subsection explains the demand side. The third subsection explains the spot-market competition. The fourth subsection shows how I solve the dynamic part of the model by backward induction.

3.1 Setup: Dynamic Discrete-Choice Game

Time is discrete with infinite horizon $t = 0, 1, 2, ..., \infty$. Two goods, old and new, are imperfect substitutes from the buyers' viewpoint. Each of these goods requires a specific generation of technology for production, old and new.

Two Types, Four States

Two types of firms ("incumbents" and "entrants") and four individual states ("old-only," "both," "new-only," and "potential entrant") exist, as illustrated in Figure 4.

Figure 4: Illustration of the Dynamic Discrete Choice

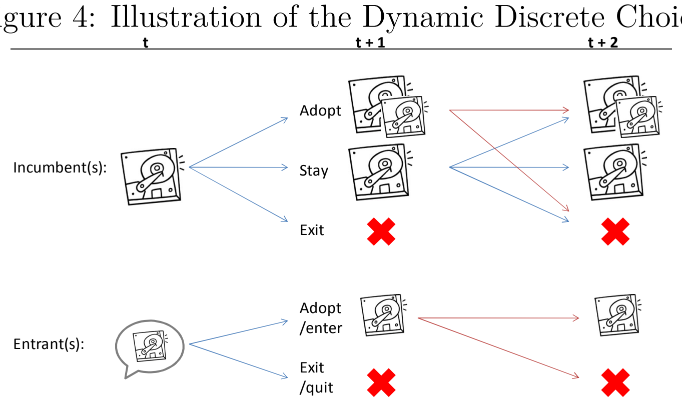

Note: "Old-only" and "New-only" firms produce 5.25- and 3.5-inch HDDs exclusively. "Both" represents incumbents that adopted the new technology, hence producing both of the two generations. "Potential entrant" is recognized (in data) by the announcement of product specifications without actual shipment.

"Incumbents" produce old goods from time 0, using the old-generation technology. In any time period, an incumbent may choose to adopt the new-generation technology by paying a sunk cost and starting to produce both old and new goods from the next period. Hence an incumbent starts in the "old-only" state (a production technology status in which a firm can produce only old goods) and may elect to transition to the “both” state (a status in which a firm can produce both old and new goods).

“Entrants” are the other type of firms. They are not active in the market at time 0. Each of them appears in a predetermined year (observed in data),[10] at which moment they may choose to adopt new technology and enter the market, or quit the prospect of entry.[11] That is, by paying a sunk cost, a potential entrant becomes an actual entrant in the subsequent period in the “new-only” state, a production technology status in which a firm can produce only new goods.[12]

Hence a firm belongs to any one of the four states, $s_{it} \in \{old, both, new, pe\}$. The industry state summarizes all firms' states, $s_t = \left(N_t^{old}, N_t^{both}, N_t^{new}, N_t^{pe}\right)$, as the numbers of firms in each of four states. Let $s_{-i,t}$ denote the numbers of competitors for firm i, which is simply s_t minus 1 in firm i's own category.

Period Profit

Firm i's single-period profit, $\pi\left(s_{it}, s_{-i,t}, W_t\right)$, depends on its own state s_{it}, its competitors' state $s_{-i,t}$, and the characteristics of demand and cost W_t. s_{it} and $s_{-i,t}$ change endogenously as a result of each firm's dynamic decision, whereas W_t evolves exogenously. Old and new

[10] I have also attempted the following alternative modeling approaches: (1) assuming a fixed number of potential entrants for all years, (2) assuming the difference between the peak number of active firms across all years (in data) and the current number of active firms as the number of potential entrants in that year, (3) assuming infinitely many potential entrants and their sequential free entry, and (4) a combination of (3) and my approach in the paper. Although these modeling approaches have their merits, I have chosen to model potential entrants in a more data-driven manner, for the following reasons. Conceptually, the primary purpose of this research is to empirically analyze incumbents' decision-making in the presence of *credible* threats of entry, rather than developing some generic model of entry. In this context, the estimation of incumbents' and entrants' innovation costs is most meaningful when we compare incumbents and the most serious contenders that are *known to exist* by the industry participants (and hence recorded in the industry publications), rather than numerous potential entrants that might be theoretically possible but practically invisible. Quantitatively my modeling choice will translate into estimating a lower bound of entrants' innovation/entry cost. The more potential entrants we assume, the higher the estimated cost of entry has to be in order to rationalize the number of actual entrants in the data.

[11] An alternative modeling approach is to allow potential entrants to wait multiple years before actual entry, but I have chosen to characterize them as ephemeral creatures, for three reasons. First, potential entrants in the data do not wait. They either enter in the following year or disappear from record. Second, no venture capitalist in Silicon Valley appears sufficiently patient to let entrepreneurs wait and burn cash for multiple years, especially in a fast-changing industry. Third, an empirical model of dynamic oligopoly needs to define the state space parsimoniously. Given my other modeling choices, potential entrants with longer life would increase the size of the state space exponentially and make its computation infeasible. I conjecture that, quantitatively, this modeling choice will affect my estimates in the following manner. On the one hand, assuming potential entrants with longer life might lower the entry cost estimate because we would not need as high an entry cost to rationalize their no entry, although how this logic might interact with the observational fact that no potential entrant waited multiple years is less obvious. On the other hand, incumbents might have stronger preemptive motives if potential entrants can wait and are hence less desperate to enter, which might alter the equilibrium strategy profile.

[12] I do not consider entrants' adoption of the old technology because it rarely happens in the data once the new technology becomes available.

goods are imperfect substitutes; that is, products are differentiated *across* generations. The logit demand system characterizes their substitution pattern (section 3.2). Products are homogeneous *within* each generation, with the spot-market competition characterized by a symmetric Cournot game (section 3.3).

Dynamic Discrete-choice Problem

Each firm aims to maximize its expected present value. The interest rate is assumed to be positive and constant over time, resulting in a constant discount factor $\beta \in (0,1)$ per period. In each period, events occur in the following order:

1. Each active firm observes its private cost shocks ε_{it}^{0}, ε_{it}^{1}, and ε_{it}^{2}, associated with exiting, staying, and innovating (if a firm's state is *old*), respectively. A potential entrant draws ε_{it}^{0} (for quitting) and ε_{it}^{2} (for entry/innovation) but not ε_{it}^{1} because it does not sell anything and hence cannot wait on the sidelines.

2. Firms make dynamic decisions, namely, exit, stay, or innovate.

3. Active firms compete in the spot market and earn profits, $\pi\left(s_{it}, s_{-i,t}, W_t\right)$.

4. Dynamic decisions are implemented. Specifically, exiting firms exit and receive their sell-off values $\phi + \varepsilon_{it}^{0}$. Staying firms receive ε_{it}^{1}. Adopting incumbents pay $\delta^t \kappa^{inc} - \varepsilon_{it}^{2}$. Potential entrants receive ε_{it}^{0} (if they quit) or pay $\delta^t \kappa^{ent} - \varepsilon_{it}^{2}$ (if they adopt/enter) to become active.

5. The industry takes on a new state, $s_{t+1} = \left(N_{t+1}^{old}, N_{t+1}^{both}, N_{t+1}^{new}, N_{t+1}^{pe}\right)$.

The decision problems for active firms in each of the three states are characterized by[13]

$$
\begin{aligned}
V_t^{old}\left(s_t\right) &= \pi_t^{old}\left(s_t\right) + \max\left\{\begin{array}{c} \phi + \varepsilon_{it}^{0}, \\ \beta E\left[V_{t+1}^{old}\left(s_{t+1}\right)|s_t\right] + \varepsilon_{it}^{1}, \\ \beta E\left[V_{t+1}^{both}\left(s_{t+1}\right)|s_t\right] - \delta^t \kappa^{inc} + \varepsilon_{it}^{2} \end{array}\right\}, \qquad (1) \\
V_t^{both}\left(s_t\right) &= \pi_t^{both}\left(s_t\right) + \max\left\{\begin{array}{c} \phi + \varepsilon_{it}^{0}, \\ \beta E\left[V_{t+1}^{both}\left(s_{t+1}\right)|s_t\right] + \varepsilon_{it}^{1} \end{array}\right\}, \text{ and} \\
V_t^{new}\left(s_t\right) &= \pi_t^{new}\left(s_t\right) + \max\left\{\begin{array}{c} \phi + \varepsilon_{it}^{0}, \\ \beta E\left[V_{t+1}^{new}\left(s_{t+1}\right)|s_t\right] + \varepsilon_{it}^{1} \end{array}\right\},
\end{aligned}
$$

[13] For notational simplicity, I suppress ε_{it}^{0}, ε_{it}^{1}, and ε_{it}^{2} from the argument of $V_t^{\cdot}\left(s_t\right)$.

subject to the perceived law of motion governing the industry state, s_t (see below). For a potential entrant, the problem is simply

$$\max \left\{ \varepsilon_{it}^{0}, \beta E \left[V_{t+1}^{new} \left(s_{t+1} \right) | s_t \right] - \delta^t \kappa^{ent} + \varepsilon_{it}^{2} \right\}.$$

Non-stationary Environment and Solution Concept

I assume the HDD industry reaches its terminal state in 1998, when market structure stops evolving. I solve the model backward over 18 years in the spirit of subgame perfect equilibrium (SPE).[14] In reality, the industry keeps evolving after the sample period, but the point is that the 5.25-inch HDDs all but disappeared by 1998, and the 3.5-inch HDDs have remained the mainstream device to date. Thus, the environment is non-stationary with this unique technological transition, along with the explosive growth of demand and the steady decline in production costs. Because I want to analyze the economic incentives surrounding the transition from the 5.25- to 3.5-inch HDDs, I believe this non-stationary setup with (effectively) a finite horizon is a more adequate representation of history during the sample period than following the more conventional practice of imposing stationarity.

Multiple Equilibria

The multiplicity of equilibria is the common cause of concern in the studies of dynamic oligopoly, which motivated the development of the two-step estimation methodologies that bypass the issue. However, the two-step approach is too data intensive for my industry/data setting (i.e., a single, non-stationary global market).

I approach this problem with simple modeling and equilibrium selection. The solution of my dynamic game boils down to solving a static, simultaneous-move entry/exit/innovation game of incomplete information for a Bayesian Nash equilibrium (BNE) for each state in each year (see section 3.4). Seim (2006) shows the uniqueness of BNE under certain market structure for a similar game. For a subset of all possible industry states, my numerical search found only one equilibrium. For example, when only one type of firms is active, symmetry allows us to pin down the BNE. Alternatively, even when multiple types of firms are in the market, a BNE appears to be unique as long as their payoffs differ significantly. However,

[14] See section 3.4 for details. Strictly speaking, I am using not SPE but perfect Bayesian equilibrium (PBE) as a solution concept. The game involves private information in the form of private cost shocks associated with dynamic actions (entry/exit and innovation). A firm never observes the realizations of these shocks for its rivals, which precludes the existence of subgames. However, the past realizations of private cost shocks do not matter *per se*, because only the current market structure (along with time and a firm's own draws of current private shocks) affects the firm's payoff. My model shares this convenient property with the Markov perfect equilibrium (MPE) models of industry dynamics. Hence our assumptions about off-path beliefs are irrelevant. For this reason, I may also proceed with sequential equilibrium instead of PBE.

for the other subset of industry states, I do encounter multiple BNE. In such cases, I select the first BNE that is reached in best-response iterations that start from the same symmetric strategy profile.[15]

I also investigated the effectiveness of an alternative "selection" approach by imposing sequential moves within each state-year, instead of simultaneous moves, which would eliminate multiplicity in the current context. However, after several months of trial and error, I found this approach infeasible and not particularly appealing for two reasons. First, computation became prohibitively expensive, because one has to find late movers' best responses to all possible actions of early movers. Computation time increased by the order of 1000, even with a relatively coarse grid for discretizing the action space. Second, this formulation embeds early-mover advantages to those classes of firms that I specify as early movers, in every state-year, which is not a desirable specification when the research question is about the incumbent-entrant heterogeneity.

Beliefs (Perceived Law of Motion)

For rules governing firms' beliefs about rivals' moves, $s_{-i,t+1}$, I assume rational expectations. That is, a firm correctly perceives how its rivals make dynamic decisions *up to private cost shocks*, $(\varepsilon_{it}^0, \varepsilon_{it}^1, \varepsilon_{it}^2)$ iid extreme value. This setup allows for dynamic strategic interactions, which are a prerequisite for incorporating cannibalization and preemptive motives into the model.[16]

With respect to the evolution of demand and production costs, I assume firms' perfect foresight. This choice reflects my analytical focus on strategic interactions and innovation costs rather than informational factors related to demand uncertainty. In my view, this assumption is simplistic but not distortionary, because firms' beliefs are homogeneous regardless of their types or individual state in any given period. Hence it is unlikely to affect the incumbent-entrant asymmetry this paper tries to explain.[17]

Model Primitives

There are dynamic and static components of model primitives. Dynamic primitives are the discount factor β, the mean sell-off value ϕ, the base sunk costs of innovation κ^{inc} and κ^{ent},

[15] I use this criterion out of practical necessity. As Pesendorfer (2010) notes, "Lyapunov stability of the best response mapping is not a convincing equilibrium selection rule in private information games." Su (2012) demonstrates the existence of other (unstable) equilibria for a similar game.

[16] I "shut down" this feature in one of the counterfactuals (no-preemption case, in section 6.4), where firms instead perceive the industry state as evolving exogenously, in the spirit of non-stationary Oblivious Equilibrium (Weintraub, Benkard, Jeziorski, and Van Roy 2008).

[17] Although learning is beyond the scope of this paper due to data limitations, adaptive expectations and information-related topics can be important, including (1) beliefs about new HDDs' profitability that are potentially heterogeneous across firms, and (2) the updating of these beliefs through own experimentation and learning from rivals.

the annual rate of sunk cost change δ, a dynamic equilibrium concept, and the informational assumptions made on firms' perceived law of motion for the industry state.

Static primitives determine the period profit function: demand parameters, cost parameters, and a static equilibrium concept. The next two subsections explain the details of these static model components.

3.2 Demand

I capture the substitution pattern across generations of HDDs using the logit model of differentiated products. The dynamic discrete-choice model (section 3.1) captures HDDs' differentiation *across* generations and assumes homogeneity *within* each generation. A buyer k purchasing an HDD of generation g enjoys utility,[18]

$$u_{kg} = \alpha_0 + \alpha_1 p_g + \alpha_2 I\left(g = new\right) + \alpha_3 x_g + \xi_g + \epsilon_{kg}, \tag{2}$$

where p_g is the price, $I\left(g = new\right)$ is the indicator of new generation, ξ_g is the unobserved characteristics (most importantly, "design popularity" among buyers, as well as other unobserved attributes such as "reliability"), and ϵ_{kg} is the idiosyncratic taste shock over generations. The outside goods offer the normalized utility $u_{k0} \equiv 0$, which represent *removable* HDDs (as opposed to *fixed* HDDs) and other storage devices.[19]

Let $\bar{u}_g \equiv \alpha_0 + \alpha_1 p_g + \alpha_2 I\left(g = new\right) + \alpha_3 x_g + \xi_g$ represent the mean utility from a generation-g HDD whose market share is $ms_g = \exp\left(\bar{u}_g\right) / \sum_l \exp\left(\bar{u}_l\right)$. The shipment quantity of generation-g HDDs is $Q_g = ms_g M$, where M is the size of the HDD market including the outside goods (removable HDDs and other storage devices). Practically, M reflects all desktop PCs to be manufactured globally in a given year. Berry's (1994) inversion provides the linear relationship,

$$\ln\left(\frac{ms_g}{ms_0}\right) = \alpha_1 p_g + \alpha_2 I\left(g = new\right) + \alpha_3 x_g + \xi_g, \tag{3}$$

where s_g represents the market share of HDDs of generation g, and ms_0 is the market share of outside goods (removable HDDs and other devices). The inverse demand is

$$p_g = \frac{1}{-\alpha_1}\left[-\ln\left(\frac{ms_g}{ms_0}\right) + \alpha_2 I\left(g = new\right) + \alpha_3 x_g + \xi_g\right]. \tag{4}$$

[18] I suppress the time subscript t for the simplicity of notation. The demand side is static in the sense that buyers make new purchasing decisions every year. This assumption is not restrictive because multi-year contracting is not common and hundreds of buyers (computer makers) are present during the sample period.

[19] Tape recorders, optical disk drives, and flash memory.

3.3 Period Competition

The spot-market competition is characterized by multi-product (i.e., old and new goods) Cournot competition.[20] Marginal costs of producing old and new goods, mc^{old} and mc^{new}, are assumed to be common across firms and constant with respect to quantity. Firm i maximizes profits

$$\pi_i = \sum_{g \in A_i} \pi_{ig} = \sum_{g \in A_i} (p_g - mc_g)\, q_{ig} \tag{5}$$

with respect to shipping quantity q_{ig} $\forall g \in A_i$, where π_{ig} is the profit of firm i in generation g, and A_i is the set of generations produced by firm i. Firm i's first-order condition with respect to its output q_{ig} is

$$p_g + \frac{\partial p_g}{\partial Q_g} q_{ig} + \frac{\partial p_h}{\partial Q_g} q_{ih} = mc_g, \tag{6}$$

with $g, h \in \{old, new\}$, $g \neq h$, if firm i produces both old and new HDDs. The third term on the left-hand side is dropped if a firm makes only one generation.

3.4 Solution of Dynamic Game by Backward Induction

I assume the state stops evolving after year T. Hence the terminal values associated with a firm's states, $s_{iT} \in \{old, both, new\}$, are

$$\left(V_T^{old}, V_T^{both}, V_T^{new}\right) = \left(\sum_{\tau=T}^{\infty} \beta^{\tau} \pi_T^{old}(s_T), \sum_{\tau=T}^{\infty} \beta^{\tau} \pi_T^{both}(s_T), \sum_{\tau=T}^{\infty} \beta^{\tau} \pi_T^{new}(s_T)\right).^{21} \tag{7}$$

In year $T-1$, an old-only firm's problem (aside from maximizing its period profit) is

$$\max\left\{\phi + \varepsilon_{i,T-1}^{0}, \beta E\left[V_T^{old}(s_T)\,|s_{T-1}\right] + \varepsilon_{i,T-1}^{1}, \beta E\left[V_T^{both}(s_T)\,|s_{T-1}\right] - \delta^{T-1}\kappa^{inc} + \varepsilon_{i,T-1}^{2}\right\}.$$

I follow Rust (1987) to exploit the property of the logit errors, $\varepsilon_{it} = (\varepsilon_{it}^0, \varepsilon_{it}^1, \varepsilon_{it}^2)$, and their (conditional) independence over time, to obtain a closed-form expression for the expected

[20] Three considerations led to the Cournot assumption. First, unlike automobiles or ready-to-eat cereals, HDD is a high-tech "commodity." Buyers chiefly consider its price and category (i.e., form factor and storage capacity), within which the room for further differentiation is limited. Second, changes in production capacity take time, and hence the spot market is characterized by price competition given installed capacities. Third, accounting records indicate that despite fierce competition with undifferentiated goods, the HDD makers seemed to enjoy non-zero (albeit razor-thin) profit margins on average.

[21] Alternatively, I may anchor the terminal values to some auxiliary data (if available) that would cover the periods after 1998, the final year of my data set. The market capitalization of the surviving firms as of 1998 would be a natural candidate, which, combined with net debt, would represent their enterprise values. However, I stopped pursuing this approach because of (1) the survivorship bias, (2) the presence of conglomerates, and (3) the omission of private firms.

value before observing ε_{it},

$$E_{\varepsilon_{i,T-1}}\left[V_{T-1}^{old}\left(s_{T-1},\varepsilon_{i,T-1}\right)|s_{T-1}\right] = \pi_{T-1}^{old}\left(s_{T-1}\right)+\gamma$$
$$+\ln\left[\exp\left(\phi\right)+\exp\left(\beta E\left[V_{T}^{old}\left(s_{T}\right)|s_{T-1}\right]\right)+\exp\left(\beta E\left[V_{T}^{both}\left(s_{T}\right)|s_{T-1}\right]-\delta^{T-1}\kappa^{inc}\right)\right],$$

where γ is the Euler constant. Similar expressions hold for the other two types:

$$E_{\varepsilon_{i,T-1}}\left[V_{T-1}^{both}\left(s_{T-1},\varepsilon_{i,T-1}\right)|s_{T-1}\right] = \pi_{T-1}^{both}\left(s_{T-1}\right)+\gamma$$
$$+\ln\left[\exp\left(\phi\right)+\exp\left(\beta E\left[V_{T}^{both}\left(s_{T}\right)|s_{T-1}\right]\right)\right],$$

and

$$E_{\varepsilon_{i,T-1}}\left[V_{T-1}^{new}\left(s_{T-1},\varepsilon_{i,T-1}\right)|s_{T-1}\right] = \pi_{T-1}^{new}\left(s_{T-1}\right)+\gamma$$
$$+\ln\left[\exp\left(\phi\right)+\exp\left(\beta E\left[V_{T}^{new}\left(s_{T}\right)|s_{T-1}\right]\right)\right].$$

In this manner, I can write the expected value functions from year T all the way back to year 0. The associated choice probabilities (policy functions) will provide a basis for the MLE (in section 4.3).

4 Estimation

My empirical approach takes three steps. First, I estimate the system of demand for differentiated products. Second, I recover the marginal costs of production implied by the demand estimates and the first-order conditions of the firms' period-profit maximization. These (static) demand and cost estimates for each year permit the calculation of period profit for each class of firms, in each year, under any market structure (industry state). Third, I embed these period-profit estimates into the dynamic discrete game of innovation and entry/exit, which I solve to estimate the sunk costs of innovation and entry/exit.

4.1 Estimation: Demand

Although the dynamic analysis will focus on the generational transition from 5.25 to 3.5-inch HDDs, the empirical demand analysis incorporates more details. In the data, the unit of observation is the combination of generation, quality, buyer category, geographical regions, and year t. For notational simplicity, I denote the generation-quality pair by "product category" j and suppress subscripts for the latter three dimensions. A buyer k purchasing an HDD of product category j, that is, a combination of generation g (diameter) and quality x (storage capacity in megabytes), enjoys utility

$$u_{kj} = \alpha_0 + \alpha_1 p_j + \alpha_2 I\left(g_j = new\right) + \alpha_3 x_j + \xi_j + \epsilon_{kj}, \tag{8}$$

where p_j is the price, ξ_j is the unobserved characteristics, and ϵ_{kj} is the idiosyncratic taste shock that is assumed iid extreme value (over buyers and generation-quality bins); that is, its cumulative distribution function is $F\left(\epsilon_{kj}\right) = \exp\left(-\exp\left(-\epsilon_{kj}\right)\right)$. Berry's (1994) inversion allows the econometrician to run a linear regression,

$$\ln\left(\frac{ms_j}{ms_0}\right) = \alpha_1 p_j + \alpha_2 I\left(g_j = new\right) + \alpha_3 x_j + \xi_j, \tag{9}$$

where ms_j represents the market share of HDDs of category j, and ms_0 is the market share of outside goods (removable HDDs). The coefficients, α_0 through α_3 and σ, are the taste parameters, which I estimate using OLS and instrumental variables (IVs).

Sources of Identification

The demand parameters are identified by the time-series and cross-sectional variations in the data (subscripts omitted for notational simplicity) as well as the logit functional form. Three dimensions of cross-sectional variation exist. First, an HDD's product category (denoted by j) is a pair of generation and quality. Two generations and 14 discrete quality levels exist, according to the industry convention. Second, data are recorded by buyer category, PC makers and distributors/end-users. Third, data are recorded by geographical category, U.S., and non-U.S.

The OLS estimation relies on the assumption that $E\left[\xi_j | p_j, g_j, x_j\right] = 0$; that is, the price of a category-j HDD is uncorrelated with that particular category's unobserved attractiveness to the buyers. However, one might suspect a positive correlation between them because an attractive product category would command both higher willingness to pay and higher cost of production.

In the IV estimation, I use the following variables as instruments for p_j: (1) the prices in the other region and user category, (2) the number of product "models" (not firms), (3) the number of years since first appearance, and (4) the time-series "innovations" in ξ_{jt}, ν_{jt}. The first IV is used by Hausman (1996) and then by Nevo (2001). The identifying assumption is that production cost shocks are correlated across markets, whereas taste shocks are not. This assumption reflects the industry context in which HDD makers from across the globe compete in both the United States and elsewhere, whereas end users of HDDs (and hence of PCs) are more isolated geographically.

The second IV was used by Bresnahan (1981) and Berry et al. (1995) and exploits the proximity of rival products (in product space), that is, the negative correlation between the number of "models," markup, and price in oligopolies. The identifying assumption is that taste shocks in any given period are not correlated with the number of "models" in a particular product category j, which are outside my model.[22]

The third IV relies on steady declines in the marginal costs of production over years. In the HDD industry, costs dropped because of design improvements, reduced costs of key components, and offshore production in Singapore, Malaysia, Thailand, and the Philippines. This overall tendency holds at the product category level as well. The identifying assumption is that taste shocks are not correlated with such time patterns on the cost side, which I assume is exogenously determined by advances in engineering.

These three IVs have been used with cross-sectional data and static competition in the literature, but their usefulness is unknown in the context of global industry dynamics. For this reason, I also investigated the results based on an alternative, time-series IV strategy in the style of Sweeting (2012), and obtained the price coefficient estimates of approximately -3.20, a range statistically indistinguishable at the 5% level from my preferred estimate of -3.28 based on the other three IVs (see section 5.1, column 4 of Table 2). This fourth IV employs an additional identifying assumption that the unobserved quality, ξ_{jt}, evolves according to an AR(1) process,

$$\xi_{jt} = \rho \xi_{jt-1} + \nu_{jt},$$

where ρ is the autoregressive parameter,[23] and ν_{jt} is the "innovation" (in the time-series sense) that is assumed iid across product categories and over time. We can form exclusion restrictions for ν_{jt} by assuming firms at t do not know the unpredictable parts ν_{jt+1} when they make dynamic decisions.[24]

[22] The following observation motivates this IV. Firms need to make "model"-introduction decisions in prior years, without observing taste shocks in particular regions/user types in the following years. More importantly, such dynamic decisions are driven by the sum of discounted present values of future profits, which is affected only negligibly by taste shocks in any particular period, regions, or user types. Hence this identifying assumption would be plausible as long as particular regions'/user types' taste shocks are not extremely serially correlated.

[23] The estimate for ρ is .41.

[24] I intend this additional IV result as a robustness check and do not use it for the subsequent analysis of dynamics, because the AR(1) assumption on the demand side may potentially introduce some conceptual inconsistency with my other assumptions on the supply-side dynamics, in which I let firms form perfect foresight about the evolution of demand (for the purpose of alleviating the computational costs). In any case, even if we use -3.20 instead of -3.28, the subsequent results do not change materially.

4.2 Estimation: Marginal Costs

For each year, we can infer the marginal costs of production, mc_{old} and mc_{new}, from equation (6), namely, the first-order conditions for the firms' static profit maximization problems. Because the unit observation in the HDD sales data is product category level–and not firm or brand level–I maintain, as identifying assumptions, symmetry across firms (up to private cost shocks) and constant marginal cost with respect to quantity.

4.3 Estimation: Sunk Costs of Innovation

I do not intend to estimate the discount factor, β, because its identification is known to be impractical in most cases (c.f., Rust 1987). Likewise, the rate of drop in sunk costs, δ, is difficult to estimate directly from the following procedure, so instead I will assume δ equals the average rate of decline in mc_{new} over time.[25]

The contribution of an old firm i in year t to the likelihood is

$$f^{old}\left(d_{it}|s_t;\phi,\kappa^{inc},\delta\right) = pr^{old}\left(d_{it}=exit\right)^{I(d_{it}=exit)} pr^{old}\left(d_{it}=stay\right)^{I(d_{it}=stay)} pr^{old}\left(d_{it}=adopt\right)^{I(d_{it}=adopt)},$$

where $pr^{old}(\cdot)$ is the probability that an old-only firm takes a particular action d_{it}:

$$\begin{aligned} pr^{old}\left(d_{it}=exit\right) &= \exp(\phi)/A, \\ pr^{old}\left(d_{it}=stay\right) &= \exp\left(E_\varepsilon V_{t+1}^{old}(s_{t+1})\right)/A, \text{ and} \\ pr^{old}\left(d_{it}=adopt\right) &= \exp\left(E_\varepsilon V_{t+1}^{both}(s_{t+1}) - \delta^t\kappa^{inc}\right)/A, \end{aligned}$$

where $A \equiv \exp(\phi) + \exp\left(E_\varepsilon V_{t+1}^{old}(s_{t+1})\right) + \exp\left(E_\varepsilon V_{t+1}^{both}(s_{t+1}) - \delta^t\kappa^{inc}\right)$. The contributions of the other three types of firms take similar forms (see Appendix).

Year t has $N_t \equiv \left(N_t^{old}, N_t^{both}, N_t^{new}, N_t^{pe}\right)$ active firms in each state, of which $X_t \equiv \left(X_t^{old}, X_t^{both}, X_t^{new}\right)$ exit and $E_t \equiv \left(E_t^{old}, E_t^{pe}\right)$ innovate. Denote the joint likelihood for year t of observing data (N_t, X_t, E_t) by $P(X_t, E_t, N_t)$. Then the overall joint likelihood for $t = 0, 1, 2, ..., T-1$ is $P(X, E, N) = \prod_{t=0}^{T-1} P(X_t, E_t, N_t)$. Thus the ML estimators for the mean sell-off value ϕ and the base sunk costs of technology adoption κ^{inc} and κ^{ent} are

$$\arg\max_{\phi,\kappa^{inc},\kappa^{ent}} \ln\left[P(X,E,N)\right].$$

[25] See section 5.2 for details and alternative assumptions.

Sources of Identification

Intuitively, I rely on a revealed-preference argument to identify the sell-off value and the sunk costs. Algorithmically, this estimation approach searches for the parameter values of $(\phi, \kappa^{inc}, \kappa^{ent})$ that best rationalize the observed choice probabilities of innovation and entry/exit in data. Thus the key source of identification is variations in entry/exit/innovation decisions over time and across the four firm classes. For example, $\hat{\phi}$ will depend on the fraction of active firms that exited (in the data), conditional on year and market structure (which determine the payoffs based on the period-profit estimates and the dynamic game model). Likewise, $\hat{\kappa}^{inc}$ and $\hat{\kappa}^{ent}$ will mostly depend on the observed fractions of innovating incumbents and potential entrants, respectively.

5 Results

This section reports the estimation results.

5.1 Results: Demand

Table 2 displays demand estimates. I employ two market definitions, broad (1 and 2) and narrow (3 and 4). The former definition aggregates observations across both regions (U.S. and non-U.S.) and user types (computer makers and distributors/end users), in a manner consistent with the industry's context of a single, global market. However, the dataset contains richer variations across regions and user types, which we can exploit for improved precision of estimates. Moreover, the Hausman-Nevo IVs become available under the narrower market definition (i.e., by region/user type).

The IV estimates in columns (2) and (4) are generally more intuitive and highly statistically significant than the OLS estimates in columns (1) and (3). The price coefficient is negative ($\hat{\alpha}_1 < 0$), whereas both smaller size (3.5-inch diameter = new generation) and quality (the log of storage capacity) confer higher benefits ($\hat{\alpha}_2 > 0, \hat{\alpha}_3 > 0$) to the buyers.

I use column (4), the logit IV estimates under the narrow market definition, as my baseline result for the subsequent analyses. I avoid using the results based on the broader market definition. Specifically, result (2) is similar to (4) and highly intuitive, but I am concerned about the limited availability of IVs and the reduced variation in data.

All four estimates incorporate year dummies and also allow for the time-varying unobserved product quality by diameter (ξ_{jt} in equations [8] and [9]; note I suppress time-subscripts in these formulae for notational simplicity). I use equation (9) to recover $\hat{\xi}_{jt}$ as residuals. Figure 5 (left panel) shows the evolution of $\hat{\xi}_{jt}$ for both old and new HDDs, which

Table 2: Logit Demand Estimates for 5.25- and 3.5-inch HDDs

Market definition:	Broad		Narrow	
Estimation method:	OLS	IV	OLS	IV
	(1)	(2)	(3)	(4)
Price ($000)	-1.66^{***}	-2.99^{***}	$-.93^{**}$	-3.28^{***}
	(.36)	(.64)	(.38)	(.58)
Diameter = 3.5-inch	$.84^{**}$	.75	1.75^{***}	$.91^{**}$
	(.39)	(.50)	(.27)	(.36)
Log Capacity (MB)	.18	$.87^{***}$	.04	1.20^{***}
	(.25)	(.31)	(.22)	(.27)
Year dummies	*Yes*	*Yes*	*Yes*	*Yes*
Region/user dummies	–	–	*Yes*	*Yes*
Adjusted R^2	.43	.29	.50	.27
Number of obs.	176	176	405	405
Partial R^2 for Price	–	.32	–	.16
P-value	–	.00	–	.00

Note: Standard errors in parentheses. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively.

I will refer to as "popularity" henceforth. These unobserved popularities of the old and new products switched in 1992, suggesting the 3.5-inch replaced the 5.25-inch as the mainstream HDD type.

Figure 5: Estimated Popularity (Unobserved Quality) and Marginal Costs

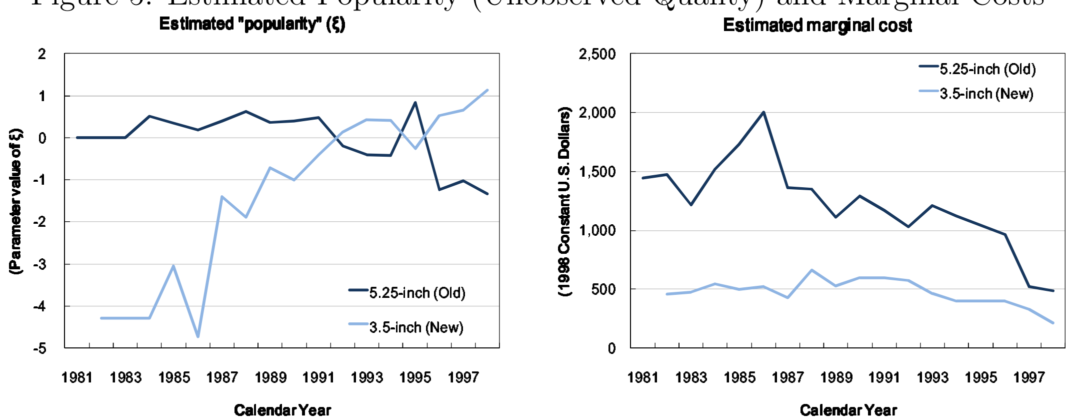

Note: Results based on the IV estimates of logit demand system with the narrow market definition (4).

5.2 Results: Marginal Costs

From the demand estimates and firms' first-order conditions, I infer marginal costs of production (Figure 5, right). The continual drop in the marginal costs reflects two tendencies in the industry. First, HDDs required increasingly fewer parts due to design improvements, reflecting advances in engineering. Second, offshore production in Singapore and other South

East Asian locations became prevalent, reducing primarily the cost of hiring engineers. Together these developments represent important channels of "process innovation."[26] The new HDDs' marginal cost declines at the average annual rate of 6.12%, which I assume equals the rate of drop in the sunk costs of adoption; that is, $\hat{\delta} = .9388$ because the adoption cost of new technology directly relates to the production of new HDDs.[27]

5.3 Results: Sunk Costs of Innovation

Table 3 shows the MLEs of the mean sell-off value, ϕ, and the base sunk costs of adopting new technology, κ^{inc} and κ^{ent}. I set $\beta = .88$ for my baseline estimates,[28] which show two important features. One is that the entrants' base sunk cost is higher than the sell-off value upon exit $\left(\hat{\phi} < \hat{\kappa}^{ent}\right)$, which implies "no free lunch;" that is, entering this market for the sole purpose of exiting and running away with the sell-off value does not pay off.

Table 3: Estimates of the Dynamic Parameters

	ML Estimates
Sell-off value (ϕ)	**4.00**
Incumbents' sunk cost $\left(\kappa^{inc}\right)$	**1.99**
Entrants' sunk cost $\left(\kappa^{ent}\right)$	**4.60**
Log likelihood	−110.2

Note: Standard errors are not available because of a non-smooth (step function-like) shape of likelihood due to strategic interactions, which precludes contraction mapping in the inner loop of NFXP and leads to either zero or large standard errors.

The other important feature is that the overall innovation/entry cost is lower for incumbents than for entrants ($\hat{\kappa}^{inc} < \hat{\kappa}^{ent}$); therefore, the seeming "inertia" of incumbents does not stem from their innate cost disadvantage. The explanation lies in other incentives, as I explore in detail with counterfactual analyses in the next section.

[26] See Igami (2013) for details on this offshoring aspect of market structure dynamics.

[27] Alternatively, one may assume increasing or time-invariant sunk costs (e.g., $\delta = 1.1$ or 1) and the results do not change qualitatively. I believe sunk costs dropping in line with production costs is more natural because both costs are concerned with the manufacturing of the same goods. In principle, one can try to estimate δ directly as a part of the dynamic model. In practice, however, such estimates tend to be unreliable, probably due to the same issue that plagues the estimation of β, the discount factor. Because changes in δ or β move almost everything in the model in the same direction, their identification seems impractical.

[28] See Appendix for more discussion and sensitivity analysis on the discount factor. The existing empirical studies of dynamic oligopolies conventionally used $\beta = .95 \sim .90$, a range that is close to my preferred model. Collard-Wexler (2011), Schmidt-Dengler (2006), Ryan (2012), and Goettler and Gordon (2011) chose $\beta = .95$, .94, .90, and .90, respectively, in their studies of the markets of concrete, hospitals, cement, and microprocessors. I believe a lower discount factor would be more adequate for the HDD industry because of its fast pace of technological changes.

Two caveats are in order regarding the interpretations of innovation/entry cost estimates. First, κ^{ent} embodies the costs of both entry and innovation for potential entrants, whereas κ^{inc} is strictly about incumbents' innovation costs. Upon entry, entrants can always elect to exit and recover ϕ, so that the truly "sunk" part of κ^{ent} is $\kappa^{ent} - \phi$. Although the exact decomposition of κ^{ent} is conceptually infeasible, I am inclined to interpret the 0.60 ($= \hat{\kappa}^{ent} - \hat{\phi}$) part of $\hat{\kappa}^{ent}$ as the pure "innovation" cost and the 4.00 ($= \hat{\phi}$) part as representing the "entry" fee. According to this interpretation, the meaning of $\hat{\kappa}^{inc} < \hat{\kappa}^{ent}$ becomes more nuanced; that is, although incumbents may enjoy overall cost advantage over entrants, the latter might be better at innovation *per se*.[29]

Second, the result $\hat{\kappa}^{inc} < \hat{\kappa}^{ent}$ does not necessarily mean incumbents are entirely free from organizational, informational, or other disadvantages. Rather, my estimates simply suggest incumbents enjoy a certain cost advantage over entrants in net terms. A possible explanation is that incumbents accumulate certain technological or marketing capabilities over the years, which outweigh other potential disadvantages associated with being larger and older.

Determining the exact contents of $\hat{\kappa}^{inc}$ and $\hat{\kappa}^{ent}$ is beyond the scope of this paper, but incumbents' overall cost advantage will have important welfare implications (see section 7).

5.4 Results: Industry State, Policy, and Value

Figure 6 (left panel) shows the evolution of the industry state in the estimated model, based on the estimated choice probabilities (policy functions). Overall, the estimated model replicates three key features of the data (Figure 3), albeit in a slightly smoother manner. First, the number of innovating incumbents (firms in "both" state) peaks at a level lower than entrants ("new-only"). Specifically, the peak numbers of "both" and "new-only" are 7 and 10 in the data, compared with approximately 5 and 10 in the model. Second, during the second half of the sample period, the survival rate is higher for innovating incumbents than for entrants, resulting in similar numbers of survivors as of 1998 (i.e., 4 "both" and 3 "new-only" in the data, compared with approximately 5 "both" and 5 "new-only" in the model). Third, the number of non-innovating incumbents ("old-only") declines precipitously during the 1980s and then more slowly during the 1990s before reaching zero in 1998. Thus the estimated model captures the key data features of innovation and market structure dynamics.

Figure 6 (right panel) shows the estimated policy function (i.e., the optimal choice probabilities) for incumbents' innovation, which starts at a rather low level, because the HDD market (and hence the profits from it) is still small. Incumbents become increasingly more

[29] I thank Lanier Benkard and Ariel Pakes for pointing to this subtlety.

Figure 6: Industry State and Innovation Policy

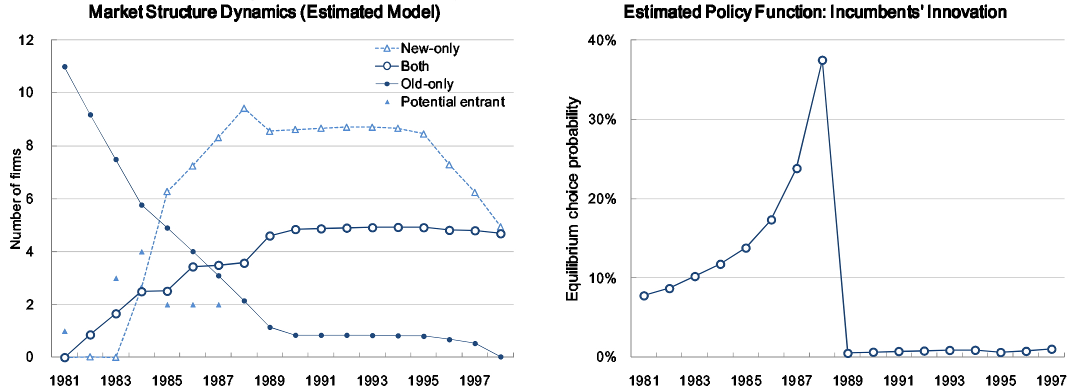

Note: Left panel displays the mean evolution of simulated industry state, based on the estimated choice probabilities (policy functions). Right panel exhibits one such sequence of policy functions, for incumbents' innovation, along the historical path of market structure.

eager to innovate in years approaching 1988, with a peak adoption rate of 37% as the PCs and HDDs become regular household goods and the new-generation HDDs gain in popularity (recall Figure 3 top left and Figure 5 left). After 1988, however, the adoption rate plummets to 1% and recovers only slightly toward 1997, the final year of dynamic decision-making.

This peak and sudden drop in the incentives to innovate epitomize preemption. Incumbents foresee the growing benefits of producing new HDDs on the eve of the demand "takeoff" and try to preempt rivals, but those who missed this timing (due to idiosyncratic cost reasons) would rather give up because, with already a dozen active firms in the production of new HDDs, innovation (and becoming the 13th new-HDD competitor) would no longer pay off.

6 "Innovator's Dilemma" Explained

This section answers the first question of the paper, namely, why incumbents are slower than entrants in innovation. I quantify the effects of the three theoretical forces that determine the incumbent-entrant timing gap in technology adoption: cannibalization, sunk-cost gap, and preemption. To measure each effect, I compare the timing gaps in the estimated baseline model with a counterfactual simulation in which that particular incentive mechanism is absent.

Figure 7 summarizes the results of the counterfactual analyses. The incumbent-entrant timing gap is measured by the percentage-point differences between incumbents' and entrants' CDF of adoption timing (c.f., Figure 2), averaged over years. "Baseline" is the estimated model's outcome (10.32 percentage points). The other three values (5.27, 16.45, and 51.73 percentage points) represent the simulated counterfactuals in which I "shut down" particular economic incentives.

Figure 7: Incumbents-Entrant Timing Gap in Innovation

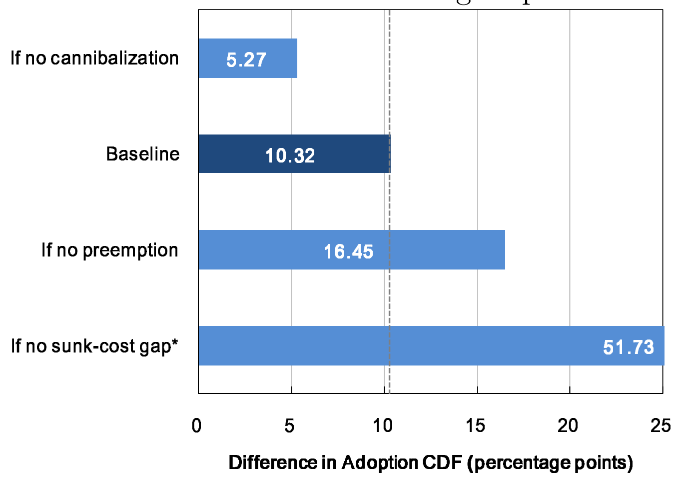

Note: * outside the graph range. Timing gap is measured by the percentage-point difference between incumbents' and entrants' CDF of adoption timing, averaged over years during the first half of the sample period. "Baseline" outcome is based on the estimated model (see previous section), whereas the other three are the counterfactual simulation results, which I explain in detail in this section.

The comparison of the counterfactuals against the baseline suggests the following: (1) cannibalization can explain 51% of the timing gap; (2) without preemptive motives (and other dynamic strategic incentives), incumbents would have further delayed innovation by as much as 59%; and (3) contrary to the prior impression of "organizational inertia," incumbents enjoy a cost advantage over entrants (and hence the elimination of this cost advantage would have led to incumbents' much longer delay). I derive these timing-gap outcomes from the simulations shown in Figure 8, which compares the evolutions of the industry state in the baseline model with three counterfactuals.

In the following three subsections, I explain the setup, result, and interpretation of each counterfactual.

6.1 Cannibalization (Counterfactual 1)

I eliminate the cannibalization concern from incumbents' optimization problem by isolating the innovation decision (production of new HDDs) from the profit maximization regarding old HDDs. I effectively split each incumbent firm into two separate entities: a "legacy" division that manufactures old HDDs and a "corporate venture" in charge of developing new HDDs. The former division acts as an independent "old-only" firm that decides each year whether to stay or exit, but without the third alternative to adopt new technology and

Figure 8: Explaining the Innovator's Dilemma (Counterfactual Series I)

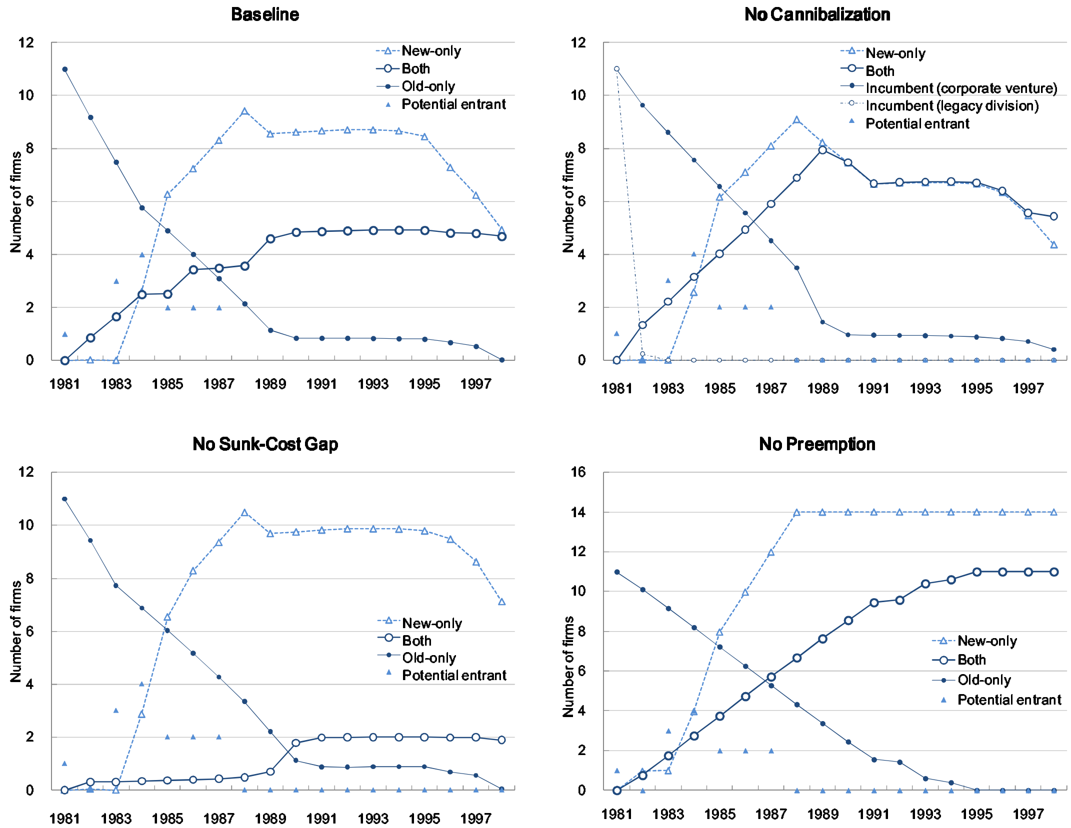

Note: "Baseline" estimated model is explained in section 5. See sections 6.1, 6.2, and 6.3 for details of the three counterfactuals.

become "both." The latter division acts like a "potential entrant" with staying power, which can choose to innovate (and become "new-only"), wait, or exit. Thus each incumbent in this no-cannibalization counterfactual is two separate firms dedicated to old and new HDDs in isolation.[30]

Incumbents (their "corporate venture" divisions, to be precise) are much more eager to adopt new HDDs than in the baseline case. Consequently, an approximately equal number of incumbents and entrants produced the new HDDs during the 1990s. Free of the cannibalization concerns regarding their own old-HDD business, more incumbents (their "corporate venture" divisions) start producing new HDDs earlier. Cannibalization can explain half of the actual timing gap between incumbents and entrants.

[30] An alternative approach to isolate the cannibalization factor would be to directly alter the HDD demand system in such a way that old and new HDDs no longer substitute for each other. Computationally, this alternative approach is easier than the approach I chose, because the latter substantially increases the effective number of firms. A drawback of the alternative approach is that the counterfactual demand system needs to be specified in a rather arbitrary manner. Two different markets must exist, of size M_t^{old} and M_t^{new}, when in fact only one market existed (with size M_t), but I see no obvious way to split M_t into M_t^{old} and M_t^{new}, and the outcomes depended heavily on the way I split the market (not reported).

6.2 Preemption (Counterfactual 2)

Preemption is a dynamic strategic motive. In an oligopolistic environment, some firms' early adoption would reduce the incremental profits available to late adopters. An incumbent has incentives to preempt other incumbents as well as potential entrants. Thus the silencing of preemption requires that firms do not perceive the evolution of industry state (the numbers of firms in "old-only," "both," and "new-only" states) as something they can influence through their own actions. In the no-preemption counterfactual, firms take the evolution of the industry as exogenous to their dynamic decisions.[31]

In the absence of preemptive motives, incumbents' delay increases substantially, to 16.45 percentage points from 10.32 percentage points in the baseline model. The number of "both" firms (adopting incumbents) grows more slowly. Each firm ignores its rivals' decisions, so the nature of the dynamic game changes fundamentally from that of strategic entry/exit to a single-agent optimal stopping problem. An incumbent does not need to act aggressively to deter rivals, so the innovation rate becomes lower and the incumbent-entrant gap wider.

6.3 Sunk Cost Gap (Counterfactual 3)

An important finding from estimating the baseline model is the sunk cost advantage of incumbents relative to entrants, the estimates of which were 1.99 and 4.60, respectively. What if incumbents no longer enjoyed this cost advantage? To eliminate the cost difference, this counterfactual sets the sunk costs at 3.30 for both incumbents and entrants ($\kappa^{inc} = \kappa^{ent} = \bar{\kappa} = 3.30$).[32]

Incumbents' innovation is discouraged. At most, only two "both" firms are active in the market. By contrast, more "new-only" firms thrive even toward the end of the 1990s. A 66% increase in sunk cost is sufficient to suppress most incumbents' innovation. Another interesting feature is the higher survival rate among entrants. Seven "new-only" firms survive until 1998, whereas only five survive in the baseline case. With few incumbents innovating and directly competing in the new HDD category, entrants can enjoy higher profits and hence improved survival prospects.

[31] One might alternatively label this counterfactual as a "no dynamic strategic interaction" or "dynamic monopolistic competition" scenario. I choose to call it a "no-preemption" case to highlight the economic incentives I believe are at the heart of firms' technology adoption decisions.

[32] Alternatively, $\bar{\kappa} = 1.99$ (i.e., entrants enjoy the low cost of incumbents) is an equally plausible setup. However, even in the baseline model, most potential entrants decide to adopt anyway, which limits the upside for entrants' adoption rate. Another possible configuration is $\bar{\kappa} = 4.60$ (i.e., incumbents are as "handicapped" as entrants). This setup results in few incumbents' adoption, only toward the end of the sample period (not reported). Although interesting, this result is too extreme to be compared with the baseline case. Therefore, I chose to show the results for the "mean" counterfactual ($\bar{\kappa} = 3.30$) for more meaningful comparison.

7 Policy Experiments

In this section, I evaluate public policies concerning innovation and competition. I conduct counterfactual simulations and compare measures of social welfare. Specifically, I experiment with two policies: (1) a broad patent on new HDDs and (2) a ban on international trade.[33] The purpose of these experiments is to inform policy design as well as to deepen our understanding of the interactions between innovation, competition, and welfare.

Table 4 summarizes the welfare analysis. Rows represent different policy simulations, including the benchmark cases. Columns list the components of social welfare: (A) consumer surplus, (B) producer surplus, (C) sell-off value upon exit, and (D) sunk costs of innovation. Social welfare is their sum. I display social welfare figures separately for the sample period (1981 through 1998) and for the years since 1999. The latter consists of the terminal values of (A) and (B) but not (C) or (D) because no more exit or innovation/entry occurs after 1998.

Table 4: Comparison of Social Welfare across Policy Experiments

(Billion $)	1981 through 1998						From 1999
	(A) Consumer Surplus	(B) Producer Surplus	(C) Sell-Off Value	(D) Innovation Cost	(A+B+C+D) Social Welfare	Change from Baseline	Social Welfare
Benchmarks							
Baseline	127.7	9.4	20.0	−39.5	117.6	0%	112.3
• U.S.	102.5	4.5	9.6	−18.0	98.6	0%	84.9
• Non U.S.	25.2	4.9	10.4	−21.6	18.9	0%	27.4
Experiments							
1. Broad Patent							
• Pre-announced	109.6	16.4	6.2	−13.9	118.2	0.5%	99.5
• Surprise in 1988	67.9	15.1	17.0	−39.0	61.0	−48.2%	30.4
2. No Trade	107.4	16.8	6.7	−31.1	99.7	−15.2%	95.5
• U.S.	88.5	12.7	0.0	−12.4	88.8	−10.0%	79.4
• Non U.S.	18.9	4.1	6.7	−18.7	11.0	−88.9%	16.2

Note: Each number is the sum of discounted present values as of 1981. * Includes government subsidies.

The largest social welfare component is consumer surplus (A). However, the second largest item, innovation costs (D), tends to vary most across different scenarios because the number and timing of innovation will drastically alter the total sunk costs. By contrast, consumer surplus stays within a narrower range as long as several firms are supplying old and new HDDs. Therefore, the final welfare outcome (A+B+C+D) mainly reflects the tradeoff between the benefits (A) and costs (D) of innovation.

[33] In Appendix, I show two theoretical benchmarks (social planner's and monopolist's solutions), as well as two additional policy experiments with (3) R&D subsidies to incumbents and (4) a ban on non-compete clauses, for a thorough assessment of the social costs and benefits of innovation.

Figure 9: Time Profile of Social Welfare (Baseline)

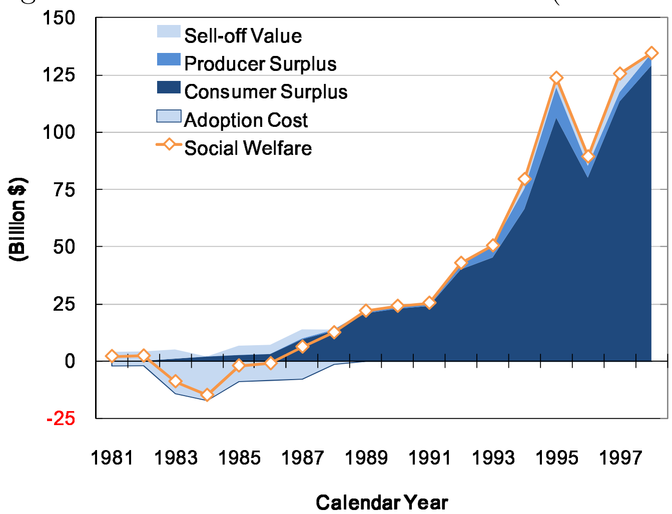

Note: Numbers are not time discounted. Excludes terminal values (i.e., surpluses accruing after 1998).

This tradeoff is clear in Figure 9, which graphs the time profile of social welfare in the baseline case. To illustrate the relative magnitude of surpluses and costs in different years, each welfare component is displayed without time discounting. Three features are noteworthy. First, the size of the consumer surplus grows over time with the growth of market size (M_t). Second, as the number of firms declines during the 1990s, the increased market power leads to higher producer surplus. The sell-off value is also visible because of continual exits. Third, adoption costs are huge and accrue in the early years, so the net social welfare is negative during most of the 1980s, when demand for HDDs is limited. Once the sunk costs are paid, however, all future surpluses count toward a net increase in social welfare. Consequently, the world enjoys a net gain of \$117.6 billion between 1981 and 1998, and \$229.9 billion if we include all years after 1998. If we consider the evolution of the HDD industry an investment project, the social "internal rate of return" would be 33.5%.[34]

7.1 Broad Patent (Experiment 1)

The question of whether broad patents encourage innovation is particularly relevant to the HDD industry, in which Rodime, a Scottish firm that was among the first adopters of the 3.5-inch technology, obtained a patent on the concept of 3.5-inch HDDs in 1986, which had previously been considered an integral part of the shared technological standards throughout

[34]The internal rate of return (IRR) is the discount rate at which the project's net present value becomes zero. The IRR of 33.5% implies the social return from the entire historical development in the HDD is positive as long as the social discount factor is above .665.

the industry (see Table 5). After years of lawsuits between Rodime and its rivals, the U.S. Court of Appeals for the Federal Circuit (CAFC) eventually rejected the claim in 1995 and 1996, but several HDD makers gave up the court battles before the rulings and agreed to pay license fees to Rodime. This event seems to be a typical "bogus patent" episode, consistent with Jaffe and Lerner's (2004) assessment of the unintended consequences of the patent system reforms during the 1980s.

Table 5: Brief History of Rodime's 3.5-inch HDD Patent Affair

Year	Events
1980	Rodime became independent from Burroughs' 5.25-inch HDD plant in Glenrothes, Scotland.
1983	Rodime became the first maker to achieve volume production of 3.5-inch HDDs.
1986	Rodime surprised the industry by obtaining a patent on the concept of a 3.5-inch drive.
	Rodime sued Miniscribe and Conner Peripheral for patent infringement.
	IBM sued Rodime, which countersued IBM.
1988	The 3.5-inch patent affair headed for a long tour of the U.S. federal court system.
	Miniscribe opted out by taking a license from Rodime.
1989	Rodime moved to Singapore for production efficiency, but neared bankruptcy and obtained some financing. The company completely overhauled top management in early 1989.
1991	Patent affair ended when IBM and Conner Peripheral, as well as Fujitsu and Alps Electric, took licenses. Several other firms were in negotiation. Rodime pursued joint ventures with Japan's JVC and firms in Taiwan and Korea, but in mid-1991 announced it would file for bankruptcy and cease manufacturing operations.
	It planned to remain active in pursuing licensing revenues from 3.5-inch HDD patents.
1994	High legal expenses and falling license revenues put financial pressure on Rodime.
1995	In September 1995, a U.S. appeals court ruled some of Rodime's patent claims invalid, a ruling in favor of Quantum. Rodime still argued other patent claims were valid, in a separate legal action against Seagate.
1996	Appeals court rulings in 1995 and 1996 appear to have weakened Rodime's negotiating position, but it continues to argue other patent claims are still valid.

Source: *DISK/TREND Reports*, various years.

Although Rodime's claims were considered outrageous in the industry at the time, and ended in a legal grey zone, studying what the welfare consequences would have been had the patent system and CAFC's rulings been different is worthwhile. I propose two separate experiments, one designed to study the *ex-ante* impact of a pre-announced broad patent regime and the other to study the impact of *ex-post* "surprise" court rulings.

In the first counterfactual, only the first innovator(s) is/are allowed to manufacture new HDDs, and the patent authority pre-announces this legal arrangement before 1981, before the game begins. The setup allows for multiple adopters as long as they belong to the first cohort.[35]

[35]The motivation for this permissive regime is to reflect the custom of defensive patenting in the computer-related industry, where rival firms tend to hold competing intellectual-property claims and engage in lawsuits and countersuits. Firms typically use patents to enhance their bargaining power in negotiating favorable terms in cross-licensing agreements. Alternatively, one might consider the patent regime in which only and

In the second counterfactual scenario, Rodime's rivals ignore the company's patent claims until 1988, when the CAFC, a centralized appellate court for patent cases established in 1982, announces its "surprise" ruling to honor Rodime's patent infringement claims, paving the way for a legal monopoly of the 3.5-inch technology.[36]

Figure 10 compares the evolutions of the industry state in the two counterfactual simulations. The pre-announced broad patent regime (left panel) induces a relatively simple market structure. Seven of the 11 incumbents decide to adopt new HDDs in 1981, after which further adoption becomes illegal. No entrants can adopt. Social welfare increases by 0.5% between 1981 and 1998, but, ironically, this gain does not come from increased innovation. Rather, social welfare increases because patents prevent innovation by many firms, thereby saving innovation costs. Broad patents function as a hard entry barrier. Because innovation cost is an important component of social welfare, limited entry translates into a huge saving in societal cost. Limited competition reduces the 1981–98 consumer surplus by 14.2% ($18.1 billion), but the innovation cost reduction of 64.8% ($25.6 billion) compensates for this loss. Seven innovators seem sufficient to ensure relatively competitive supply and pricing. The innovation timing is front-loaded to 1981, which is a favorable development from the viewpoint of pro-innovation policy.

Figure 10: Policy Experiment 1 (Ex-ante and Ex-post Effects of Broad Patent)

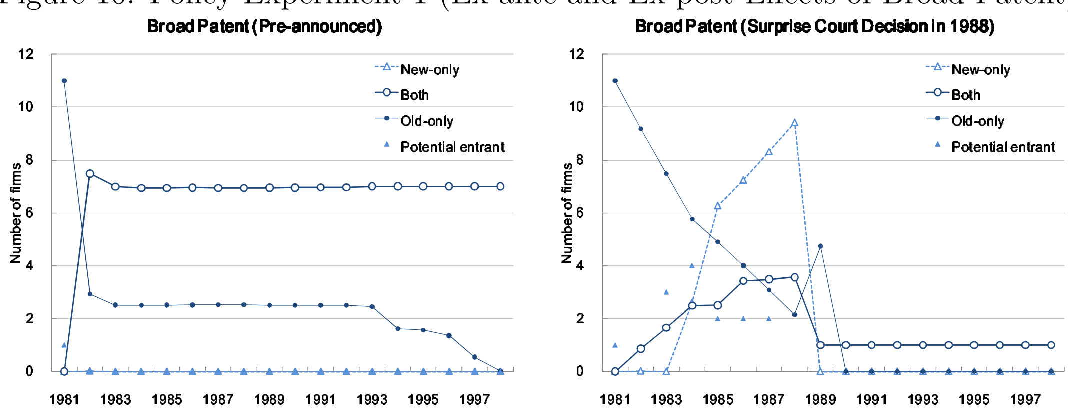

Note: For details of counterfactuals, see section 7.1

In the second, "surprise court ruling" experiment (right panel), the industry evolves just as in the baseline case until 1988, the year when the CAFC hypothetically entrenched Rodime's legal monopoly. I assume that due to the preliminary injunction to stop infringing activities, all "new-only" firms immediately go out of business, and all "both" firms except

exactly one firm can make 3.5-inch HDDs. I omit this scenario because the outcome is almost equivalent to the benchmark "monopolist" case.

[36]Patent lawsuits usually end with the payment of damages. For simplicity, this experiment assumes the payments are large enough to make Rodime's rivals indifferent between staying versus exiting.

Rodime are forced back to "old-only" status by 1989, a position that is so unattractive the firms all subsequently elect to exit the industry. In stark contrast with the previous experiment, social welfare drops by 48.2% in the ex-post patent scenario because most firms had already paid the sunk costs and started production of 3.5-inch HDDs by 1988, so that no cost saving exists as in the "pre-announced" patent regime. Instead, the only major change is that the industry becomes a monopoly from 1990 and consumers thus suffer.

In the existing literature, Lerner (1994) found a positive correlation between broader patents and increased innovative activities, whereas Sakakibara and Branstetter (2001) found no measurable relationship. By contrast, the results of my experiments point to the possibility that allowing for broad patents might decrease innovative activities. These diverging results are not necessarily at odds with each other, for two reasons. First, even when the number of total innovators decreases in the long run, if the timing is earlier, the number of innovators will be higher in the short run. In other words, a tradeoff seems to exist between short- and long-run levels of innovative activities, and patents may change this balance. These changes in the timing and number of innovations embody the ex-ante effect of patents. Second, because patents are legal monopoly rights by definition, if one firm's patent on a particular product/technology is strictly honored and protected, rival firms must be excluded from the same product/technology. The second experiment highlights this ex-post effect of broad patents.

7.2 No International Trade (Experiment 2)

Disputes over intellectual property rights often become international trade disputes in the computer-related industries. For example, Apple sued HTC, a leading manufacturer of the Android-based smart phone in Taiwan, in July 2011 for two patent infringement claims related to touch-screen technologies in iPhones. HTC countersued Apple, but, meanwhile, the U.S. Commerce Department started contemplating banning U.S. imports of HTC's smart phones. A similar dispute between Apple and Sumsung led to a temporary ban on Samsung's Galaxy tablet computer in Germany and Australia. Given the global nature of competition in these high-tech industries, intellectual property and competition policy issues tend to become international trade issues at the same time. Although the international trade of HDDs avoided political interventions during the sample period, other sectors of the computer industry were full of trade disputes.[37] Because the HDD industry data are relatively free from political complications, they provide a clean laboratory in which to experiment with

[37]For example, Japan's Ministry of International Trade and Commerce restricted the imports and manufacturing of IBM's computers during the 1960s. The U.S. Commerce Department imposed "anti-dumping" duties on Samsung's memory chips (DRAM) in the early 1990s.

the welfare impact of trade barriers.

Figure 11: Policy Experiment 2 (No International Trade)

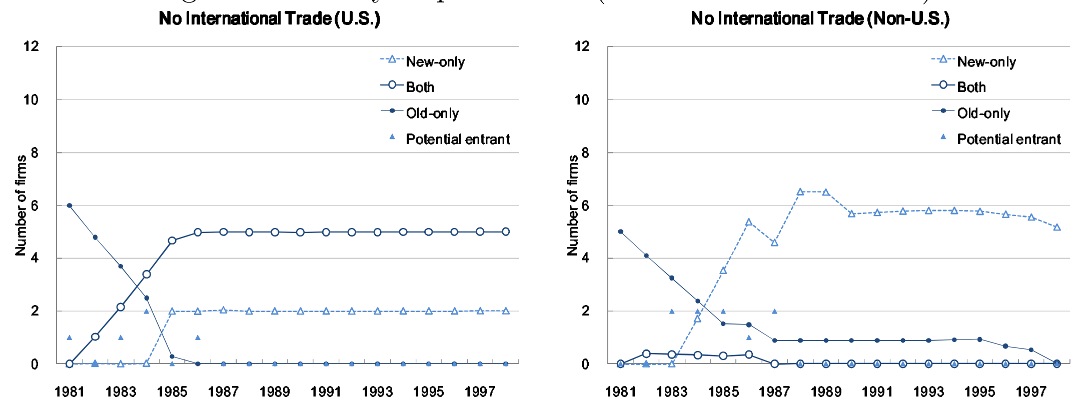

Note: For details of counterfactuals, see section 7.4.

My no-trade counterfactual examines an environment in which no HDD trade occurs between the United States and the rest of the world. On the demand side, the world market (its pool of potential buyers) is split into the U.S. and non-U.S. On the supply side, American firms can only serve the U.S. market, whereas non-American firms serve the rest of the world. Thus this experiment is one of alternative market sizes and market structures.

In the United States and elsewhere, consumer surpluses as well as the total innovation costs decrease. The net result is a 15.2% decrease in worldwide social welfare. Interestingly, the effect is asymmetric between the two regions. U.S. social welfare drops by only 10% while the rest of the world suffers an 88.9% reduction.

On the supply side, trade restrictions are anti-competitive in limiting the numbers of effective competitors both within and outside the United States. Moreover, the smaller market size on the demand side implies smaller spoils of innovation for potential adopters, discouraging the production of new HDDs. These two effects lead to more concentrated market structures. The latter effect translates into a lower innovation rate and hence cost savings. Because these cost savings accrue to producers, the effect of trade barriers is positive for U.S. producers. However, non-U.S. producers' surplus declines because the number of firms originating from outside the United States is disproportionately high compared to the size of the market. Thus, whereas the concentration effect (i.e., a smaller number of firms compete) dominates in the United States, to the advantage of American HDD makers, outside the United States, the market size effect (i.e., the existence of fewer buyers of HDDs) dominates. In other words, consumers across the globe would suffer from the lack of trade, and so would non-American producers. American firms would be the only parties to benefit from the ban on trade.

8 Conclusion

This paper presents an industry equilibrium model of innovation and entry/exit with dynamic strategic interactions, which I estimate using an 18-year panel of HDD market data. Contrary to Schumpeter's (1934) earlier conjecture, incumbents actually enjoy an advantage over entrants in terms of innovation sunk costs—a result that echoes his later observation (Schumpeter 1942). This finding implies incumbents are not slower than entrants because of an efficiency handicap.

Then what explains the incumbent-entrant timing gap? The first set of counterfactual simulations quantifies the three theoretical forces that determine incumbents' delayed adoption timing compared to entrants. The results suggest cannibalization between old and new technologies/products can explain at least a half of the dilemma. Because old and new HDDs substitute for each other, the introduction of new HDDs would dampen profits from the old HDDs. Taking this "replacement effect" into consideration, rational incumbents can expect relatively small benefits from innovation compared to new entrants. Thus creative destruction (i.e., the process of transition from old to new technologies along with firms' turnover) takes place not necessarily because of old firms' lack of creativity, but because of their rational unwillingness to destroy old sources of profits. This finding explains why new entrants, despite their cost disadvantage, sometimes manage to creatively destroy old winners in the industry: the *rational* innovator's dilemma.

Another finding from the first set of simulations is that dynamic strategic incentives, such as preemption, are important determinants of innovation and industry evolution. Without strategic interactions, the incumbent-entrant timing gap would have been wider by 59%. Thus certain aspects of industry dynamics call for explicit—although often computationally burdensome—modeling of game-theoretic interactions among firms.

In the second set of simulations, I experiment with typical public policies related to innovation and competition. The main finding is that competition is generally pro-innovation, as highlighted in the counterfactual simulation of a "no international trade" scenario. However, the welfare implication of competition is more nuanced. When many firms adopt new technology, the duplication of effort often turns out to be socially wasteful. Three factors explain this outcome. First, the estimated sunk cost of innovation is economically significant. Second, once more than four or five firms compete in the same (new) product category, a further increase in the number of adopters adds little to consumer surplus. Third, although buyers generally do benefit from early production of new goods, the market size is small during the early years. Thus the earlier innovation by many firms would result in negligible gains at sizeable costs. This finding explains why the simulated R&D subsidies lead to more innovation but reduced social welfare, and also why a (pre-announced) broad-based patent, which represents a fairly anti-competitive intellectual property regime, results in increased social welfare.

Overall, the historical welfare performance of the HDD industry in the data seems difficult to beat, even with carefully designed policy interventions, so Schumpeter (1942) was probably right when he observed the world economy was benefiting from an approximately sufficient amount of competition and innovation. Although some of these results would be specific to the HDD market, the economic incentives studied here are general and could be expected to operate in many oligopolistic markets. Similar trajectories of creative destruction are also widely observed across industries, especially in computer-related sectors. These considerations lead me to believe this paper's analytical framework, as well as its empirical findings, could be applicable elsewhere: old winners can survive the "gale of creative destruction" only through creative *self*-destruction.

Appendix

A.1 Regressing Innovation Timing on Firm Characteristics

As a preliminary, descriptive analysis, I regress the timing of innovation (i.e., each firm's year of initial shipment of new-generation HDDs) on firm characteristics. Table 6 reports the estimation results based on a standard duration model (Cox proportional hazard estimates). The estimates suggest incumbents are 50% less likely than entrants to innovate in a given year, even after controlling for all of the observed characteristics.

Table 6: Preliminary Regression of Innovation Timing on Firm Characteristics

	Duration model (Cox proportional hazard estimates)					
Dependent variable: Decision to Innovate	(1)		(2)		(3)	
Firm characteristics						
Incumbent	$.41^{***}$	(.06)	$.53^{***}$	(.09)	$.50^{***}$	(.09)
Initial generation of entry						
5.25-inch	–	(–)	.84	(.18)	.71	(.19)
3.5-inch	–	(–)	$.52^{***}$	(.12)	$.46^{***}$	(.12)
2.5-inch	–	(–)	$.34^{***}$	(.10)	.91	(.40)
Organizational Form						
Specialized HDD start-up	–	(–)	1.04	(.22)	1.01	(.22)
Computer maker ("backward" integration)	–	(–)	1.39	(.29)	1.32	(.28)
HDD component maker ("forward" integration)	–	(–)	.55	(.21)	$.54^{*}$	(.20)
Region of Origin						
U.S.	–	(–)	1.04	(.32)	1.18	(.37)
Asia	–	(–)	1.91^{**}	(.58)	1.90^{**}	(.58)
Europe (West)	–	(–)	1.01	(.37)	1.20	(.45)
Europe (East)	–	(–)	$.14^{*}$	(.15)	$.16^{*}$	(.17)
Industry state						
Number of firms	–	(–)	–	(–)	1.16^{***}	(.05)
Number of firms squared	–	(–)	–	(–)	$.99^{***}$	(.00)
Number of subjects (firms)	437		437		437	
Number of innovations	190		190		190	
Time at risk of innovation	2,591		2,591		2,591	
Number of observations	1,842		1,842		1,842	
Log likelihood	−1,018		−997		−990	

Note: Coefficients greater (less) than 1 indicate higher (lower) propensities to innovate. Omitted categories are "Potential entrant," "8-inch," "Other electronics maker (horizontal diversification)," and "Brazil." ***, **, and * indicate signifiance at the 1%, 5%, and 10% levels, respectively. Standard errors in parentheses.

I employ two different definitions of entrants in these regressions. The first definition is data-driven and narrower: a potential entrant is recognized when a new firm announces the product specifications without actually manufacturing or shipping them. The second definition is more conceptually motivated and broader: for all of the potential entrants that announced or shipped HDDs, I count as "time at risk (of innovation)" all years since the industry-wide establishment of new-generation standards. That is, regardless of whether these new firms were already incorporated or not, I pretend that their founders were considering the innovation and entry, from the moment the industry consensus emerged on the physical size of new-generation HDDs.

Econometrically, regressions based on the first definition would over-estimate the entrants' propensity to innovate because this definition recognizes entrants only when they were most serious about innovation. In contrast, regressions based on the second definition would under-estimate entrants' propensity to innovate because all years until actual shipments are interpreted as evidence against their innovativeness, regardless of whether they existed as firms or not. Table 6 reports only the results based on the second definition, in order to show that entrants were twice more likely to innovate than incumbents, even when I use the definition that would bias against such findings.

A.2 Likelihood Function

The contributions of the other three types of firms are

$$\begin{aligned} f^{both}\left(d_{it}|s_t;\phi\right) &= pr^{both}\left(d_{it}=exit\right)^{I(d_{it}=exit)}\,pr^{both}\left(d_{it}=stay\right)^{I(d_{it}=stay)}, \\ f^{new}\left(d_{it}|s_t;\phi\right) &= pr^{new}\left(d_{it}=exit\right)^{I(d_{it}=exit)}\,pr^{new}\left(d_{it}=stay\right)^{I(d_{it}=stay)}, \text{ and} \\ f^{pe}\left(d_{it}|s_t;\kappa^{ent},\delta\right) &= pr^{pe}\left(d_{it}=quit\right)^{I(d_{it}=quit)}\,pr^{pe}\left(d_{it}=adopt\right)^{I(d_{it}=adopt)}, \end{aligned}$$

where

$$\begin{aligned} pr^{both}\left(d_{it}=exit\right) &= \frac{\exp\left(\phi\right)}{\exp\left(\phi\right)+\exp\left(E_\varepsilon V^{both}_{t+1}\left(s_{t+1}\right)\right)}, \\ pr^{both}\left(d_{it}=stay\right) &= \frac{\exp\left(E_\varepsilon V^{both}_{t+1}\left(s_{t+1}\right)\right)}{\exp\left(\phi\right)+\exp\left(E_\varepsilon V^{both}_{t+1}\left(s_{t+1}\right)\right)}, \\ pr^{new}\left(d_{it}=exit\right) &= \frac{\exp\left(\phi\right)}{\exp\left(\phi\right)+\exp\left(E_\varepsilon V^{new}_{t+1}\left(s_{t+1}\right)\right)}, \\ pr^{new}\left(d_{it}=stay\right) &= \frac{\exp\left(E_\varepsilon V^{new}_{t+1}\left(s_{t+1}\right)\right)}{\exp\left(\phi\right)+\exp\left(E_\varepsilon V^{new}_{t+1}\left(s_{t+1}\right)\right)}, \\ pr^{pe}\left(d_{it}=quit\right) &= \frac{\exp\left(0\right)}{\exp\left(0\right)+\exp\left(E_\varepsilon V^{new}_{t+1}\left(s_{t+1}\right)-\delta^t\kappa^{ent}\right)}, \text{ and} \\ pr^{pe}\left(d_{it}=adopt\right) &= \frac{\exp\left(E_\varepsilon V^{new}_{t+1}\left(s_{t+1}\right)-\delta^t\kappa^{ent}\right)}{\exp\left(0\right)+\exp\left(E_\varepsilon V^{new}_{t+1}\left(s_{t+1}\right)-\delta^t\kappa^{ent}\right)}. \end{aligned}$$

The joint likelihood for year t of observing data (N_t, X_t, E_t) is

$$
\begin{aligned}
P\left(X_t, E_t, N_t\right) = & \binom{N_t^{old}}{X_t^{old}}\binom{N_t^{old}-X_t^{old}}{E_t^{old}} pr^{old}\left(d_{it}=exit\right)^{X_t^{old}} \\
& \times pr^{old}\left(d_{it}=stay\right)^{N_t^{old}-X_t^{old}-E_t^{old}} pr^{old}\left(d_{it}=adopt\right)^{E_t^{old}} \\
& \times \binom{N_t^{both}}{X_t^{both}} pr^{both}\left(d_{it}=exit\right)^{X_t^{both}} pr^{both}\left(d_{it}=stay\right)^{N_t^{both}-X_t^{both}} \\
& \times \binom{N_t^{new}}{X_t^{new}} pr^{new}\left(d_{it}=exit\right)^{X_t^{new}} pr^{new}\left(d_{it}=stay\right)^{N_t^{new}-X_t^{new}} \\
& \times \binom{N_t^{pe}}{E_t^{pe}} pr^{pe}\left(d_{it}=adopt\right)^{E_t^{pe}} pr^{pe}\left(d_{it}=quit\right)^{N_t^{pe}-E_t^{pe}}
\end{aligned}
$$

A.3 Sensitivity Analysis on Discount Factor

Each column represents a set of coefficient estimates given a particular value of the discount factor, β.

Table 7: Estimates of the Dynamic Parameters

(Billion \$)	$\beta=.94$	$\beta=.92$	$\beta=.90$	$\beta=.88$	$\beta=.86$	$\beta=.84$	$\beta=.82$
Sell-off value (ϕ)	12.51	8.87	5.69	**4.00**	2.74	2.04	1.41
Incumbents' sunk cost $\left(\kappa^{inc}\right)$	2.62	2.50	2.38	**1.99**	1.89	1.77	1.72
Entrants' sunk cost $\left(\kappa^{ent}\right)$	8.69	8.25	7.08	**4.60**	2.95	2.06	1.42
Log likelihood	−129.3	−117.1	−113.3	−110.2	−107.3	−104.2	−101.6

Note: Values of β outside this range (between .82 and .94) result in either computational errors or unintuitive parameter estimates (e.g., negative ϕ). Standard errors are not available because of a non-smooth (step function-like) shape of likelihood due to strategic interactions, which precludes contraction mapping in the inner loop of NFXP and leads to either zero or very large standard errors.

Estimates tend to decrease with β because a lower discount factor implies a lesser value of doing business in the HDD market, hence lower values associated with entry/exit, too. The log likelihood also increases with β, suggesting a slightly better fit (in terms of the choice probabilities, upon which the likelihood is based) for estimates based on lower β. As β approaches .80, however, the sell-off value estimate $\hat{\phi}$ drops close to zero and then turns negative (not reported). The adoption cost estimates $\hat{\kappa}^{inc}$ and $\hat{\kappa}^{ent}$ also turn negative with lower βs. Such results are economically implausible because, taken literally, they would suggest firms are somehow "penalized" upon exit (instead of earning sell-off value) and "rewarded" upon innovation (instead of paying sunk cost).

Consequently, I choose $\beta = .88$ estimates as my baseline result: the choice that best reflects the tradeoff between data fit and economic sensibility. On the one hand, the $\beta = .88$

estimates fit data considerably better than $\beta = .94$, .92, or .90. On the other hand, the coefficient estimates with $\beta = .88$ are sufficiently higher than zero, allowing straightforward interpretations. Furthermore, the simulation of market structure based on $\beta = .88$ performs better than simulations based on lower β in terms of the peak number of firms in "both" and "new-only" states, that is, active players that have adopted new technology, because a lower β leads to lower values of doing business. This lowering of value induces too many exits and too few adoptions to match the number of firms observed in the data. Hence a higher β is desirable in this aspect of data fit (in state evolution, in contrast to the fit in choice probabilities as measured by log likelihood in Table 3).

A.4 Other Counterfactual Simulations

This section presents additional counterfactual simulations to better understand the nature of welfare tradeoffs between the costs and benefits of innovation.

Theoretical Benchmarks

To find theoretical benchmarks, I calculate the welfare profiles under the hypothetical scenarios in which a social planner and a monopolist, respectively, make production decisions. In both cases, only one incumbent adopts new technology because that is the cheapest way to produce new HDDs. The planner implements marginal-cost pricing ($p_g = mc_g$), whereas the monopolist charges profit-maximizing prices. Interestingly, both the planner and the monopolist will choose to innovate in the same year, 1986, when the overall demand for HDD just begins to take off. Their optimal innovation timing coincides because the time profiles of their costs are identical, as they both operate a single incumbent firm and consider its optimal adoption year. Their benefits calculations also have similar time profiles based on the overall size of potential demand, although the ways in which the planner and the monopolist evaluate (appropriate) surpluses are different, namely, consumer surplus versus producer surplus.

R&D Subsidy (Experiment 3)

Is subsidizing the adoption of new HDDs socially worthwhile? R&D subsidies are rampant, as arguing against the virtues of innovation is difficult. Whether subsidies are really desirable is another issue, which I explore in this experiment.

I set $\kappa^{inc} = 0$ in this counterfactual, letting the government subsidize 100% of incumbents' innovation cost. This setup does not mean innovations are socially costless; rather, the cost simply accrues to the public sector, which is reflected in Table 4 as adoption cost (D). This R&D subsidy specifically targets incumbents because their costs are lower. More fundamentally, for the government to identify and contact potential entrants is not always

Table 8: Comparison of Social Welfare across Policy Experiments

(Billion $)	1981 through 1998						From 1999
	(A) Consumer Surplus	(B) Producer Surplus	(C) Sell-Off Value	(D) Adoption Cost	(A+B+C+D) Social Welfare	Change from Baseline	Social Welfare
Benchmarks							
1. Baseline	127.7	9.4	20.0	−39.5	117.6	0%	112.3
• U.S.	102.5	4.5	9.6	−18.0	98.6	0%	84.9
• Non U.S.	25.2	4.9	10.4	−21.6	18.9	0%	27.4
2. Planner	410.7	0.0	0.0	−0.8	409.9	248.7%	485.0
3. Monopolist	60.3	16.6	0.0	−0.8	76.0	−35.3%	30.4
Experiments							
1. Broad Patent							
• Pre-announced	109.6	16.4	6.2	−13.9	118.2	0.5%	99.5
• Surprise in 1988	67.9	15.1	17.0	−39.0	61.0	−48.2%	30.4
2. No Trade	107.4	16.8	6.7	−31.1	99.7	−15.2%	95.5
• U.S.	88.5	12.7	0.0	−12.4	88.8	−10.0%	79.4
• Non U.S.	18.9	4.1	6.7	−18.7	11.0	−88.9%	16.2
3. R&D Subsidy	127.2	9.4	20.7	−45.4*	111.9	−4.8%	104.0
4. Ban NCC	129.6	7.6	18.5	−47.5	108.2	−8.0%	112.5

Note: Each number is the sum of discounted present values as of 1981. * Includes government subsidies.

feasible.[38]

Subsidies accelerate incumbents' adoption. The peak number of innovating incumbents ("both") reaches eight in 1987, compared with the baseline result of five in 1989. However, the total social welfare is 4.8% lower than in the baseline case because more innovators means higher total adoption costs.

Figure 12: Policy Experiments 3 and 4

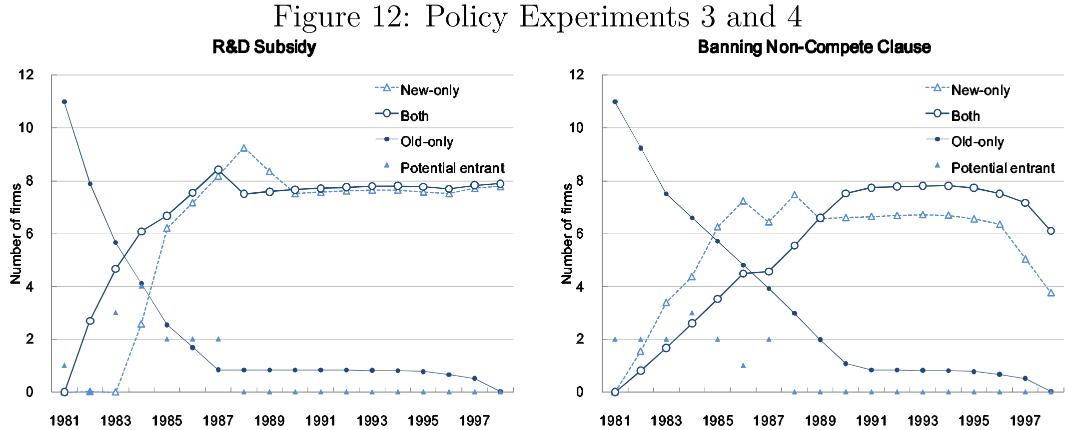

Note: See Appendix for details of counterfactuals.

Subsidies successfully encourage more incumbents to adopt, and in earlier years. However, this aggressive adoption does not translate into a significant increase in consumer surplus

[38]Raising enough capital to lobby the government is also difficult for potential entrants.

for two reasons. First, a higher number of exits offsets a higher number of adopters. Second, potential entrants are less likely to enter due to preemption, and even if they do, they are more likely to exit.

This result echoes some of the findings in the existing research. The effect of R&D subsidy has been an active area of innovation studies, where a crucial policy question is whether R&D subsidies increase or decrease innovative activities in general, and private R&D spending in particular.[39] My experiment suggests R&D subsidy would increase the number of adopters, but public R&D spending would substitute for private R&D spending without inducing a complementary increase in private R&D spending, because of the once-and-for-all nature of innovation in the technological context of the HDD industry. Although the nature of R&D efforts varies by industry, this experiment suggests a potential explanation for why studies tend to find small effects of public R&D spending in inducing private R&D.

Banning Non-compete Clauses (Experiment 4)

As Scherer (1980) and Bresnahan (2003) emphasized as the (then) chief economists of the Federal Trade Commission, entrants are important sources of innovation.[40] Therefore, entry facilitation is a potential channel through which public policy may promote innovation. One possible policy instrument concerns the legality of the private contracts governing the ease of starting new businesses. Specifically, the non-compete clause (NCC, or the covenants not to compete [CNC]) is the type of employment contract that limits the actions of personnel working at an existing firm and considering starting their own businesses (or working for rival firms). Many HDD makers are headquartered in the American state of California, where NCCs are banned. Nevertheless, exceptions to this ban exist, such as the cases involving mergers, and the practice was common during the sample period.[41] Of the 14 potential entrants in my data, five Californian firms are founded by former managers and engineers from existing HDD makers, a typical Silicon Valley phenomenon. Thus a strict ban on NCC may have important welfare implications through innovation and competition.

In this counterfactual, the five potential entrants founded by industry veterans appear two years earlier than in the actual data.

The two-year front-loading of the five entrants results in their earlier adoption. This front-loading does not materially alter the trajectory of incumbents. Social welfare decreases by 8% because of increased adoption costs.

As in the case of more aggressive adoption by incumbents (section 7.2), earlier adoption

[39] See Klette et al.'s (2000) survey.

[40] Managers seem to agree, too. See Grove (1996) for the former Intel CEO's account.

[41] For instance, Finis Conner had to wait for two years before leaving Seagate Technology to found Conner Technology in the late 1990s, because he had been a party to the merger between Seagate Technology and Conner Peripheral in years prior to his departure.

tends to generate limited extra consumer surplus while raising the adoption costs substantially. Consumer surplus does not increase substantially, because the market size (the number of potential buyers) is still small during the early years. Adoption costs increase because they are both initially high and discounted less.

References

[1] Acemoglu, Daron, and Dan Cao. 2010. "Innovation by Entrants and Incumbents." Manuscript, MIT.

[2] Adner, Ron, and Peter Zemsky. 2005. "Disruptive Technologies and the Emergence of Competition." *RAND Journal of Economics*, 36(2): 229-254.

[3] Arrow, Kenneth J. 1962. "Economic Welfare and the Allocation of Resources to Invention." In R.R. Nelson (Ed.), *The Rate and Direction of Economic Activity.* N.Y., Princeton University Press.

[4] Benkard, C. Lanier. 2004. "A Dynamic Analysis of the Market for Wide-Bodied Commercial Aircraft." *Review of Economic Studies*, 71: 581–611.

[5] Berry, Steven T.. 1994. "Estimating Discrete-Choice Models of Product Differentiation." *RAND Journal of Economics*, 25(2): 242-262.

[6] ———, James Levinsohn, and Ariel Pakes. 1995. "Automobile Prices in Market Equilibrium." *Econometrica*, 841–90.

[7] Bresnahan, Timothy F. 1981. "Departures from Marginal-Cost Pricing in the American Automobile Industry." *Journal of Econometrics*, 17: 201–27.

[8] ———. 2003. "Pro-Innovation Competition Policy: Microsoft and Beyond." CPRC Discussion Paper, Competition Policy Research Center, Fair Trade Commission of Japan.

[9] Chesbrough, H. 1999. "Arrested development: The experience of European hard disk drive firms in comparison to U.S. and Japanese firms." *Journal of Evolutional Economics*, 9: 287–330.

[10] Christensen, Clayton M.. 1993. "The Rigid Disk Drive Industry: A History of Commercial and Technological Turbulence." *Business History Review*, 67: 531-588.

[11] ———. 1997. *The Innovator's Dilemma.* New York, HarperBusiness.

[12] Cohen, W.M. 2010. "Fifty Years of Empirical Studies of Innovative Activity and Performance." Chapter 4 of B. Hall and M. Rosenberg eds., *Handbook of the Economics of Innovation*, Volume 1. Amsterdam: Elsevier.

[13] Collard-Wexler, Allan. 2011. "Mergers and Sunk Costs: An application to the ready-mix concrete industry." *Econometrica*, forthcoming.

[14] Ericson, Richard, and Ariel Pakes. 1995. “Markov-Perfect Industry Dynamics: A Framework for Empirical Work.” *Review of Economic Studies*, 62: 53-82.

[15] Franco, April Mitchell, and Darren Filson. 2006. “Spin-Outs: Knowledge Diffusion through Employee Mobility.” *RAND Journal of Economics*, 37(4): 841–60.

[16] Fudenberg, Drew, and Jean Tirole. 1986. *Dynamic Models of Oligopoly*. London, Routledge.

[17] Gilbert, Richard. 2006. “Looking for Mr. Schumpeter: Where Are We in the Competition-Innovation Debate?” NBER.

[18] ———, and David Newbery. 1982. “Preemptive Patenting and the Persistence of Monopoly.” *American Economic Review*, 72(2): 514-526.

[19] Goettler, Ronald, and Brett Gordon. 2011. “Does AMD spur Intel to innovate more?” *Journal of Political Economy*, 119(6): 1141-1200.

[20] Grove, Andrew S.. 1996. *Only the Paranoid Survive: How to Exploit the Crisis Points that Challenge Every Company and Career*. New York, Currency Doubleday.

[21] Hausman, Jerry. 1996. “Valuation of New Goods under Perfect and Imperfect Competition.” In T. Bresnahan and R. Gordon, eds., *The Economics of New Goods*, Studies in Income and Wealth, Vol. 58, Chicago: National Bureau of Economic Research.

[22] Henderson, Rebecca M. 1993. “Underinvestment and Incompetence as Responses to Radical Innovation: Evidence from the Photolithographic Alignment Industry.” *RAND Journal of Economics*, 24(2): 248-270.

[23] ———, and Kim B. Clark. 1990. “Architectural Innovation: The Reconfiguration of Existing Product Technologies and the Failure of Established Firms.” *Administrative Science Quarterly*, 35: 9–30.

[24] Hopenhayn, Hugo A.. 1992. “Exit, Selection, and the Value of Firms.” *Journal of Economic Dynamics and Control*, 16(3-4): 621-653.

[25] Igami, Mitsuru. 2013. “Offshoring as Process Innovation.” Manuscript, Yale University.

[26] Jaffe, Adam B., and Josh Lerner. 2004. *Innovation and Its Discontents: How Our Broken Patent System Is Endangering Innovation and Progress, and What to Do about It*. Princeton, NJ: Princeton University Press.

[27] Jovanovic, Boyan. 1982. "Selection and the Evolution of Industry." *Econometrica*, 50(3): 649-670.

[28] King, A.A., and C.L. Tucci. 2002. "Incumbent entry into new market niches: The role of experience and managerial choice in the creation of dynamic capabilities." *Management Science*, 48: 171–86.

[29] Klepper, Steven. 1996. "Entry, Exit, Growth, and Innovation over the Product Life Cycle." *American Economic Review*, 86(3): 562-583.

[30] ———, and Peter Thompson. 2006. "Submarkets and the Evolution of Market Structure." *RAND Journal of Economics*, 37(4): 861-886.

[31] Klette, Tor Jakob, and Samuel Kortum. 2004. "Innovating firms and aggregate innovation." *Journal of Political Economy*, 112(5): 986–1018.

[32] Klette, Tor J., Jarle Moen and Zvi Griliches. 2000. "Do Subsidies to Commercial R&D Reduce Market Failure? Microeconometric Evaluation Studies." *Research Policy*, 29(4–5): 471–495.

[33] Lerner, Joshua. 1994. "The Importance of Patent Scope: An Empirical Analysis." *RAND Journal of Economics*, 25: 319–333.

[34] Lerner, Josh. 1997. "An Empirical Exploration of a Technology Race." *RAND Journal of Economics*, 28(2): 228-247.

[35] McKendrick, David G., Richard F. Doner, and Stephan Haggard. 2000. *From Silicon Valley to Singapore: Location and Competitive Advantage in the Hard Disk Drive Industry*. Stanford, CA: Stanford University Press.

[36] Nelson, Richard R., and Sidney G. Winter. 1978. "Forces generating and limiting concentration under Schumpeterian competition." *Bell Journal of Economics*, 9: 524–48.

[37] ———. 1982. "The Schumpeterian tradeoff revisited." *American Economic Review*, 72: 114–32.

[38] Nevo, Aviv. 2001. "Measuring Market Power in the Ready-to-eat Cereal Industry." *Econometrica*, 69(2): 307–42.

[39] Pesendorfer, Martin. 2010. "Estimation of (Dynamic) Games: A Discussion." Manuscript prepared for the World Congress (Econometric Society Monographs) on Advances in Economics and Econometrics: Theory and Applications.

[40] Reinganum, Jennifer F.. 1983. “Uncertain Innovation and the Persistence of Monopoly.” *American Economic Review*, 73(4): 741–748.

[41] Rust, John. 1987. “Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher.” *Econometrica*, 55(5): 999-1033.

[42] Ryan, Stephen P.. 2012. “The Costs of Environmental Regulation in a Concentrated Industry.” *Econometrica*, 80(3): 1019–61.

[43] Sah, R.J. and J.E. Stiglitz. 1986. “The architecture of economic systems: Hierarchies and polyarchies.” *American Economic Review*, 76: 716–27.

[44] Sakakibara, Mariko and Lee Branstetter. 2001. “Do Stronger Patents Induce More Innovation? Evidence from the 1988 Japanese Patent Law Reforms.” *RAND Journal of Economics*, 32: 77–100.

[45] Scherer, F.M.. 1980. *Industrial market structure and economic performance*. Rand MacNally, Chicago, IL.

[46] ——— and D. Ross. 1990. *Industrial Market Structure and Economic Performance*. Houghton Mifflin, Boston, MA.

[47] Schmidt-Dengler, Philipp. 2006. “The Timing of New Technology Adoption: The Case of MRI.” Manuscript, London School of Economics.

[48] Schumpeter, Joseph A. 1934. *The Theory of Economic Development*. Cambridge, MA: Harvard University Press.

[49] ———. 1942. *Capitalism, Socialism and Democracy*. New York: Harper.

[50] Seim, Katja. 2006. “An Empirical Model of Firm Entry with Endogenous Product-Type Choices.” *RAND Journal of Economics*, 37(3), 619-640.

[51] Stoneman, Paul, and Giuliana Battisti. 2010. “The Diffusion of New Technology.” Chapter 17 of B. Hall and M. Rosenberg eds., *Handbook of the Economics of Innovation*, Volume 1, Amsterdam: Elsevier.

[52] Su, Che-Lin. 2012. “Estimating Discrete-Choice Games of Incomplete Information: A Simple Static Example.” Manuscript, Chicago Booth.

[53] Sutton, John. 1998. *Technology and market structure*. MIT Press, Cambridge, MA.

[54] Sweeting, Andrew. 2012. "Dynamic Product Positioning in Differentiated Product Markets: The Effect of Fees for Musical Performance Rights on the Commercial Radio Industry." Manuscript, Duke University.

[55] Tushman, Michael L., and Philip Anderson. 1986. "Technological Discontinuities and Organizational Environments." *Administrative Science Quarterly*, 31(3): 439-465.

[56] Weintraub, Gabriel Y., C. Lanier Benkard, Przemyslaw Jeziorski, and Benjamin Van Roy. 2008. "Nonstationary Oblivious Equilibrium." Manuscript, Columbia University.