Estimating the Innovator's Dilemma: Structural Analysis of Creative Destruction in the Hard Disk Drive Industry, 1981–1998

Mitsuru Igami

Yale University

This paper studies strategic industry dynamics of creative destruction in which firms and technologies experience turnover. Theories predict that cannibalization between existing and new products delays incumbents' innovation, whereas preemptive motives accelerate it. Incumbents' cost (dis)advantage relative to that of entrants would further reinforce these tendencies. To empirically assess these three forces, I estimate a dynamic oligopoly model using a unique panel data set of hard disk drive manufacturers. The results suggest that despite strong preemptive motives and a substantial cost advantage over entrants, cannibalization makes incumbents reluctant to innovate, which can explain at least 57 percent of the incumbent-entrant innovation gap.

I. Introduction

Technologies come and go, taking generations of companies with them. Empirical studies have shown that new ventures and smaller firms ac-

I thank my dissertation advisers at University of California Los Angeles: Daniel Ackerberg, Hugo Hopenhayn, Edward Leamer, Mariko Sakakibara, Connan Snider, and Raphael Thomadsen. For suggestions, I thank the editor, Ali Hortaçsu, the two anonymous referees, Ron Adner, Lanier Benkard, Steven Berry, Allan Collard-Wexler, Michael Dickstein, Pinelopi Goldberg, Jinyong Hahn, Philip Haile, Heli Koski, Phillip Leslie, Rosa Matzkin, Matthew Mitchell, Ichiro Obara, Taisuke Otsu, Ariel Pakes, Martin Pesendorfer, Marc Rysman, Che-Lin Su, John Sutton, Kosuke Uetake, Yong Hyeon Yang, and Mark Zbaracki, as well as seminar participants at the International Industrial Organization Conference 2011, Trans-Atlantic Doctoral Conference 2011, Roundtable for Engineering Entrepreneurship Research 2011, American Economic Association 2013, US Patent and Trademark Office–Searle Center Conference on Innovation Economics 2014, Western Ontario Ivey, University of British

Electronically published May 1, 2017
[*Journal of Political Economy*, 2017, vol. 125, no. 3]

count for a large share of innovations,[1] and anecdotal evidence suggests that old winners tend to lag behind entrants even when introducing a new technology is not too difficult.[2] Thus, to understand the simultaneous turnovers of firms and technologies (i.e., creative destruction), we need to study why incumbent firms would appear either reluctant about or incapable of making drastic innovations.[3] Who innovates better and survives longer is a fundamental question for economists and a vital question for businesses. Moreover, the welfare consequence of public policies hinges critically on the subtle trade-off between the costs and benefits of innovation (e.g., Bresnahan 2003). For these purposes, this paper presents a structural empirical analysis of creative destruction, focusing on the technological transition from the 5.25- to 3.5-inch generations in the hard disk drive (HDD) industry, in which only about half of all incumbents ever innovated into the 3.5-inch generation.

The incumbent-entrant innovation gap has been the subject of many studies, including Christensen's (1993) doctoral dissertation on the history of the HDD industry, which he later extended to multiple industries in a best-selling business book entitled *The Innovator's Dilemma* (1997). Despite the ease with which casual empiricists talk about the phenomena,

Columbia Sauder, Harvard (Economics and Harvard Business School Strategy), Yale (Economics and School of Management Marketing), Dartmouth Tuck, Notre Dame, Maryland, University of California Berkeley (Haas: Business and Public Policy), New York University Stern, Kyoto, Keio, Columbia, Wisconsin–Madison, Stanford Graduate School of Business (Economics and Marketing), Princeton, Chicago Booth (Marketing), Hitotsubashi, London School of Economics, and Connecticut. I thank Clayton Christensen for encouragement, 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 Center for International Business Education and Research, and the Dissertation Year Fellowship is gratefully acknowledged. Data are provided as supplementary material online

[1] See Scherer (1965), Gellman Research Associates (1976, 1982), the Futures Group (1984), Pavit, Robson, and Townsend (1987), and Acs and Audretsch (1988). See Cohen (2010) for a survey of the empirical literature on firm characteristics and innovation.

[2] For example, Apple's smartphones came, and Nokia's feature phones went. Amazon sells everything from electronic books to disposable diapers, whereas Borders liquidated its bookshops. These examples may seem extreme, but old winners tend to lag behind new entrants even when introducing a new technology is not too difficult. Blockbuster started its online video-streaming business with thousands of paying subscribers as early as 2006, when Netflix was a mere DVD mailing service. Likewise, Eastman Kodak developed its own digital cameras long before the advent of digital photography but did not commercialize new technologies fast enough. These examples suggest that, even when an incumbent becomes the first adopter of a new technology, it may not have the incentives to make sufficient investments to become a dominant firm in the new product category. The intensive margin of investment is not explicitly modeled in this paper but could play an important role in some industries.

[3] This sentence follows Arrow's (1962) definition of drastic innovation as a technological change that is sufficiently large to alter the existing market structure. What exactly constitutes "drastic," "radical," or "disruptive" innovation is ambiguous in most studies, with the notable exceptions of Arrow (1962), Tushman and Anderson (1986), Henderson and Clark (1990), Henderson (1993), Ehrnberg and Sjöberg (1995), Christensen (1997), and Tripsas (1997).

objective measurement of this gap is not a trivial task. By definition, *potential* entrants do not appear in historical records until they become *actual* entrants; hence, we cannot observe unsuccessful potential entrants in usual data sets. This censoring problem makes the incumbent-entrant difference an elusive concept for empirical studies. With this caveat in mind, let us look at figure 1, which shows two different measures of what Christensen and others may have meant. The top panel shows a gap in the eventual numbers of innovators, in terms of the cumulative numbers of major firms that started shipping 3.5-inch HDDs, among incumbents (i.e., the manufacturers of 5.25-inch HDDs with measurable market shares) and entrants (i.e., those who entered the HDD market for the first time with 3.5-inch HDDs and attained measurable market shares), respectively. The bottom panel expresses similar numbers in terms of the fractions of all firms that could have innovated. That is, for incumbents, the denominator is the number of all 5.25-inch HDD makers in 1981, including fringe firms with negligible market shares. For entrants, the denominator is the number of all entrants that announced their intent to manufacture and ship 3.5-inch HDDs at some point in time between 1981 and 1998.[4] This measure of potential entrants is better than a simple count of actual entrants but is still imperfect because the number and the timing of such announcements are likely to be endogenous (i.e., influenced by the underlying demand and technological conditions, as well as the competitive environment). My empirical analysis will fully account for this problem.[5] Regardless of how we measure (potential) entrants, however, figure 1 shows an important fact: only a fraction of all incumbents ever innovated into the 3.5-inch generation. Thus, why incumbents delay innovation remains a valid question, and I will estimate a model of this industry to quantify their economic incentives.

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, organizational inertia may result in higher costs of innovation for incumbents. Economic theory as well as case studies suggest that as firms grow larger and older, their R&D efficiency

[4] In the HDD industry, serious start-ups typically announce their product specifications at an early stage of development before raising additional capital from venture funds, for the purpose of assessing potential customers' interests. My data source records such announcements. See online app. A.1.1 for further details.

[5] See Sec. V.C for the details of how I specify entry and its implications for estimation.

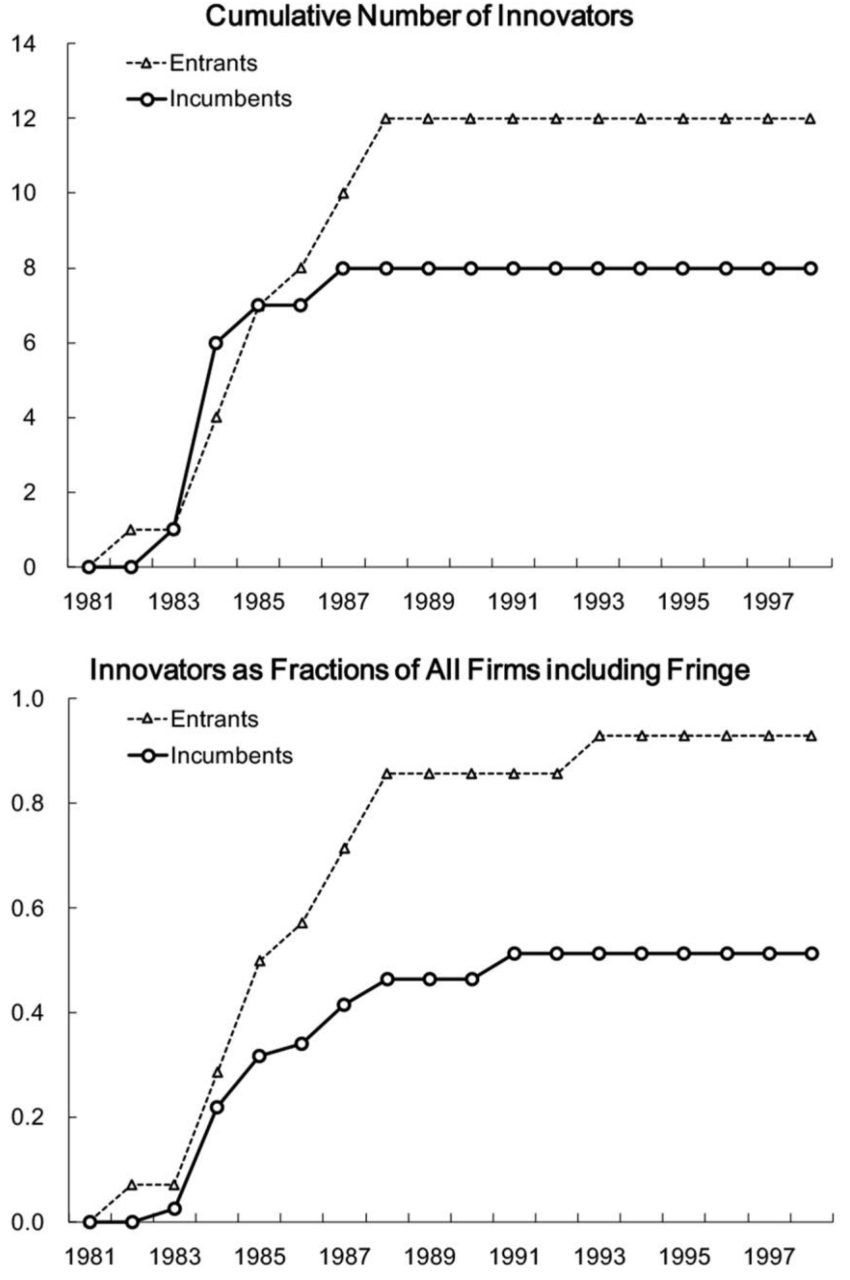

FIG. 1.—The incumbent-entrant innovation gap. The top panel plots the timing of the initial shipment of 3.5-inch HDDs separately for incumbents (i.e., firms already active in the 5.25-inch generation) and entrants (i.e., firms that appeared for the first time as producers of 3.5-inch HDDs). The bottom panel expresses similar numbers in terms of the fractions of all firms that could have innovated. See the text for details.

diminishes, although, a priori, hypothesizing that incumbency confers some advantages due to accumulated R&D capital is equally plausible.[6] Hence, whether incumbents have a cost advantage or disadvantage is an open empirical question. Third, market structure dynamics play an important, countervailing role, because theories predict that 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 market size. 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, estimate it using the data from the HDD industry, one of the best-known examples of creative destruction in which generations of firms and technologies turned over (fig. 2), and conduct counterfactual simulations to assess the empirical importance of each force.[7]

My data set consists of two parts. The first part records the industry-average price and aggregate shipment quantity for each category of HDDs, where a product category is defined as a combination of form-factor generation (e.g., 5.25- and 3.5-inch) and quality in terms of information storage capacity (e.g., 100 megabytes, 500 megabytes, and 1 gigabyte). The second part is a panel of the world's HDD manufacturers that contains information on their entry, exit, and production status (e.g., whether each firm is actively shipping 5.25- and/or 3.5-inch HDDs).

I use these data along with a simple structural model to quantify these economic forces in four steps. First, the data on the aggregate prices and quantities allow me to estimate a discrete-choice (logit) demand model. The estimated substitution pattern between the old and new HDDs will determine the extent of potential cannibalization when an incumbent firm decides to produce both products. Second, I assume homogeneous Cournot competition within each generation of HDDs, the first-order conditions of which imply the variable costs of (and profits from) manufacturing the old and new HDDs. This oligopolistic environment incorporates preemption motives, because the higher the number of firms selling new HDDs, the less profit each manufacturer earns, and because each firm forms rational expectations about its rivals' decisions in the dy-

[6] The existing literature suggests various reasons for incumbents' inertia, such as bureaucratization (Schumpeter 1934), information screening (Arrow 1974), hierarchy (Sah and Stiglitz 1986), loss of managerial control (Scherer and Ross 1990), and cognitive or relationship reasons (Grove 1996; Christensen 1997).

[7] See Christensen (1993, 1997), Lerner (1997), Chesbrough (1999), McKendrick, Doner, and Haggard (2000), King and Tucci (2002), and Franco and Filson (2006).

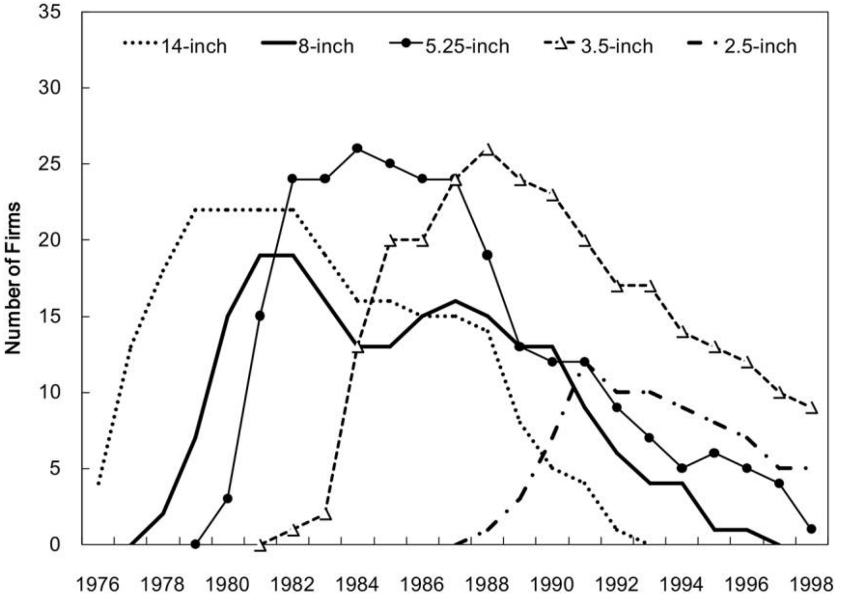

FIG. 2.—Shifting generations of technology. Multigeneration firms are counted multiple times. Mergers and acquisitions were not a major channel of exit during the sample period, but their importance grew in the 2000s. See Igami and Uetake (2015) for a dynamic structural analysis of endogenous mergers and innovation in this industry.

namic part of my model. Third, I embed these implied period profits into a dynamic oligopoly game of entry, exit, and innovation and estimate its key parameters (the firms' fixed costs of operation, as well as the incumbents' and entrants' sunk costs of innovation) using the panel data of firms' entry/exit and production status.

I model the firms' investment problems as a discrete choice between exiting, staying, and innovating (in the case of incumbents) or between entering/innovating and not (in the case of potential entrants), with alternating moves among different types of firms and private cost shocks associated with each of the dynamic discrete alternatives (assumed independent and identically distributed [iid] extreme value across firms and alternatives). For each candidate vector of parameters, I solve this dynamic game for a perfect Bayesian equilibrium (PBE) by backward induction, construct the likelihood of observing the actual choices in the data, and obtain as a maximum likelihood estimate the parameter vector that best rationalizes the observed entry/exit/innovation patterns. In the fourth step of my empirical analysis, I use the estimated model to conduct counterfactual simulations in which particular theoretical forces are absent and compare the resulting industry dynamics with those from the baseline model.

The results suggest that incumbents are reluctant innovators because of cannibalization, which can explain 57 percent of the incumbent-entrant innovation gap in the technological transition from 5.25-inch to 3.5-inch HDDs, despite strong preemptive motives and an advantage over entrants in innovation efficiency.[8] In other words, incumbents' rational reluctance rather than intrinsic inability caused the observed delay of their innovations. This finding resonates with the statement by Finis Conner, the cofounder of Seagate Technology and the founder of Conner Peripherals, that Seagate was not willing to invest in the 3.5-inch technology "because it encroached on their 5.25-inch business."[9] Moreover, the finding that incumbents actually enjoyed superior R&D efficiency (i.e., lower sunk costs of innovation) highlights the importance of distinguishing between the positivistic and normative sides of the process of creative destruction. An empirical observation that generations of firms and technologies experience turnover simultaneously does not automatically translate into the social desirability of entrants' innovations, because a hypothetical social planner would rather use an incumbent firm to produce new goods in the current context. Thus, measuring reluctance and inability separately is crucial for a sound welfare judgment.

In terms of public policy implications, counterfactual experiments imply that an idealized patent system may improve social welfare by as much as 63 percent if it worked perfectly as an ex ante incentive scheme. However, such a policy is probably infeasible in a complex technology space. A more realistic ex post granting of monopoly rights exhibits disappointing welfare performance (−91 percent), and so does the more nuanced patent regime with license fees (−5.8 percent to +3.3 percent). The failure of these typical "pro-innovation" government interventions might appear to be negative findings, but they also imply that the actual history of the HDD industry performed rather well in the subtle trade-off between competition and innovation. This finding resonates with Schumpeter's (1942) conjecture that the process of creative destruction ensures competition and innovation in the long run.

I have organized the rest of the paper as follows. The following paragraphs discuss how this paper relates to the existing literature. Section II explains why the technological transition from 5.25-inch to 3.5-inch HDDs provides an ideal empirical context in which to study creative destruction and motivates the subsequent modeling assumptions with descriptive data analysis. Section III describes the model. Sections IV and V explain the estimation procedure and results. Section VI quantifies

[8] This innovation gap is measured by the eventual difference in the cumulative numbers of innovators, as shown in fig. 1 (top).

[9] From the author's interview with Finis Conner in Corona del Mar, California, on April 20, 2015. See online app. A.1.0

the three economic forces. Section VII discusses two alternative hypotheses in the spirit of Christensen's (1997) cognitive and organizational biases and Jovanovic's (1982) learning and selection. Section VIII evaluates welfare performances of hypothetical patent policies. Section IX presents conclusions.

The online appendix contains supplementary materials including interviews and preliminary regressions (A.1), some details about the likelihood function (A.2), additional sensitivity analyses (A.3), additional counterfactual simulations including a hypothetical international intellectual property dispute (A.4), and methodological details of the full-solution approach to estimate a dynamic game in a nonstationary global market (A.5).

Related literature.—This paper studies the process of creative destruction in which firms and technologies experience turnover. Although commonly understood as a turnover of technologies alone, Schumpeter's (1942) original characterization centered on the existential threat that innovations pose to established firms, as well as their procompetitive effects on prices and quantities in the long run.[10] A key ingredient to the simultaneous turnover of firms and technologies is the existing firms' slow responses to new technologies and competitors, which is why this paper highlights the strategic industry dynamics of creative destruction, aiming to contribute to a large literature on competition and innovation (see Gilbert [2006] and Cohen [2010] for surveys).

My methods build on two strands of the empirical industrial organization literature, that is, investment and entry/exit. The most closely related papers are Benkard (2004), Schmidt-Dengler (2006), and Goettler and Gordon (2011), each of which estimates a dynamic oligopoly game of innovation using a full-solution approach. By contrast, Kim (2013) and Hashmi and Van Biesebroeck (2016) employ Bajari, Benkard, and Levin's (2007) two-step approach. Whereas these papers analyze innovation decisions of a few incumbent firms without market entry or exit, my model needs to incorporate both incumbents and entrants (up to more than two dozen firms) because I focus on the incumbent-entrant heterogeneity in innovation incentives.[11]

[10] Schumpeter (1942) asks why "the modern standard of life of the masses evolved during the period of relatively unfettered 'big business'" between 1870 and 1930, to which he subsequently answers, "the competition from the new commodity, the new technology, the new source of supply, the new type of organization—competition which commands a decisive cost or quality advantage and which strikes not at the margins of the profits and the outputs of the existing firms but at their foundations and their very lives . . . will in the long run enforce behavior very similar to the perfectly competitive pattern" (chaps. 5, 7).

[11] Xu (2008) and Aw, Roberts, and Xu (2009) study empirical settings with many firms by applying monopolistic competition frameworks, which tend to mute strategic interactions between incumbents and entrants.

Thus entry/exit is the second strand of related literature. Thematically, this paper shares a focus on the long-run evolution of market structure with the literature on industry dynamics.[12] Methodologically, I build on Seim's (2006) static entry game with incomplete information, as well as on the study by Aguirregabiria and Mira (2007), who study a dynamic entry game with incomplete information.[13] My modeling and estimation approaches diverge from the more conventional framework within a stationary environment and an infinite horizon. Specifically, I embed Seim-style discrete choice with private information into a nonstationary, finite-horizon, sequential-move dynamic game; focus on type-symmetric strategies to avoid multiple equilibria; and use Rust's (1987) nested fixed-point estimation approach.

Creative destruction has also been studied in the growth and management literature.[14] Macroeconomic models have typically abstracted from strategic interactions among incumbents and entrants, whereas the management literature has provided thick qualitative description. This paper aims to provide a microeconomic middle ground with a structural analysis of the famous episode in business history.

II. Industry and Data

This section explains why the HDD industry is particularly suitable for the study of innovation and industry evolution and describes its key features to motivate my subsequent modeling choices. Online appendix A.1 explains further details with more descriptive analysis.

A. HDD: Canonical Case of Creative Destruction

The HDD industry provides a particularly fruitful example for the study of technological change and industry dynamics, because it is the canonical example of "disruptive innovation" (Christensen 1993, 1997). 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 fig. 2). To facilitate the coordination of formats across various computer-related industries, HDD makers shared

[12] See Gort and Klepper (1982), Jovanovic (1982), Klepper and Graddy (1990), Sutton (1991, 1998, 2013), Hopenhayn (1992), Klepper (1996, 2002), and Klepper and Simons (2000).

[13] Other related papers include Fudenberg and Tirole (1986), Pakes, Ostrovsky, and Berry (2007), Pesendorfer and Schmidt-Dengler (2008), Ryan (2012), Egesdal, Lai, and Su (2014), and Su (2014). See the online appendix for further methodological considerations.

[14] Examples in the growth literature include Klette and Kortum (2004), Lentz and Mortensen (2008), and Acemoglu and Cao (2010). Examples in the management literature include Tushman and Anderson (1986), Henderson and Clark (1990), and Henderson (1993), among others.

the technological road map and the key concepts of new-generation products including the diameter of disks, but the actual commercialization process was totally up to individual firms' efforts. The introduction of new HDDs 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 (see online app. A.1.0 for further technological details).

Along with each generation, a cohort of firms came and went, many of which delayed the adoption of a newer technology. In each of the four transitions (across five generations), only about half of all incumbents (i.e., firms already active in the previous generation) ever innovated into a new generation. Even among those that did, their timing was approximately 2 years later than that of entrants (i.e., firms that appeared for the first time as producers of new-generation HDDs). Those that never innovated gradually disappeared along with the 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. 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.

B. Data

I manually construct a comprehensive panel of the world's HDD manufacturers from *DISK/TREND Reports* (1977–99), an annual publication series edited by the HDD experts in Silicon Valley.[15] I digitize 1,378 firm-year observations, each of 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. An auxiliary data set, also from *DISK/TREND Reports*, contains the aggregate prices and shipment quantities of HDDs. For each year, the reports record the average transaction price and total quantity for each of the generation-quality categories (five generations and 14 quality levels in total).

I analyze the technological transition from the 5.25- to 3.5-inch generations, which I will henceforth call the "old" and "new" generations. This subsample of the data set spans 18 years (1981–98) and 259 firm-years. I concentrate on these generations because they competed directly with

[15] Researchers have repeatedly confirmed the accuracy, relevance, and comprehensiveness of the record. See Christensen (1993, 1997), Lerner (1997), McKendrick et al. (2000), and Franco and Filson (2006).

each other in the desktop personal computer (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.[16] These two generations were historically the most important of all generations in terms of shipment volume and revenue (see fig. A3 in online app. A.1.5).[17]

The data set records each firm's production status over 23 years, along with characteristics such as their technological generation with which they entered the HDD market for the first time, organizational form (specialized, vertically integrated, or horizontally diversified), and region of origin. Table 1 shows that less than 30 percent of the firms that produced 14-, 8-, or 5.25-inch HDDs ever moved on to produce 3.5-inch HDDs, which draws our attention to these incumbents' innovation incentives. By contrast, other firm characteristics do not appear to covary with innovation timing in a statistically significant manner. This observation is confirmed by preliminary regressions using a duration model (see app. A.1.1, table A1), with a possible exception of Asian firms, and hence I choose to abstract from what appear to be secondary dimensions of firm heterogeneity in modeling a dynamic game.

DISK/TREND Reports record HDD sales by product category at the aggregate level, and not at the firm or brand level (table 2), which is an important data consideration in designing an estimable model.[18] Consequently, firm-level market shares are not recorded either, which precludes the identification of a model with rich heterogeneity. For these reasons, my modeling efforts focus on the incumbent-entrant heterogeneity, and not the firm-level heterogeneity within each class of firms.

The two generations of HDD experienced a fast growth in volume and a steady decline in price (fig. 3, top panel). The average quality (information storage capacity) of HDDs improved at an approximately constant rate (left-bottom panel). These developments were typical of those in

[16] For example, Digital Equipment Corporation (DEC) produced both minicomputers and 8-inch HDDs for these machines but subsequently failed to survive in the age of PCs and 5.25-inch HDDs. Because DEC was primarily a minicomputer manufacturer that happened to backward-integrate the 8-inch HDD manufacturing processes, its fate would be better understood in the context of the broader computer industry (i.e., the transition from minicomputers to PCs) than as a matter of competition and innovation in the HDD market.

[17] For these reasons, Christensen's (1993) historical study also devotes most attention to the transition from 5.25- to 3.5-inch HDDs.

[18] James Porter, late editor of *DISK/TREND Reports*, explained that this limitation was due to his confidentiality agreements with the HDD firms. Nevertheless, the industry participants saw high commercial value in the reports, because the average price and aggregate output data by product category were sufficiently informative in the market of high-tech "commodity" goods that are characterized by little brand differentiation.

TABLE 1
DESCRIPTIVE STATISTICS: FIRM CHARACTERISTICS AND INNOVATION TIMING

Firm Characteristics	Number of Firms	Ever Produced 3.5-Inch: Count	Ever Produced 3.5-Inch: Fraction (%)	Initial 3.5-Inch Production: Mean Year	Initial 3.5-Inch Production: Standard Deviation
Initial diameter of entry:					
14-inch	41	10	24.4	1985.5	2.3
8-inch	21	3	14.3	1987.3	3.2
5.25-inch	66	19	28.8	1986.6	2.5
3.5-inch	36	31	86.1[a]	1987.3	3.3
Other	14	0	.0	. . .	. . .
Organizational form:					
Specialized HDD start-up	74	24	32.4	1986.9	3.5
Computer maker	52	19	36.5	1986.6	2.3
HDD component maker	13	2	15.4	1991.0	5.7
Other electronics maker	39	17	43.4	1986.4	2.3
Region of origin:					
United States	101	22	21.8	1986.2	2.6
Asia	46	30	65.2	1987.1	2.9
Europe (west)	18	4	22.2	1986.8	4.9
Europe (east)	3	0	.0	. . .	. . .
Brazil	10	6	60.0	1987.0	3.3
Total	178	62	34.8	1986.8	2.9

SOURCE.—*DISK/TREND Reports*.
[a] Less than 100% of the firms that announced intentions to produce 3.5-inch HDDs actually did.

many computer-related industries. The right-bottom panel shows the numbers of firms in three of the four technological states: (1) "old only," (2) "both," (3) "new only," and (4) "potential entrant." Incumbents start in state 1 and transit to state 2 after innovation. Entrants start in state 4 and transit to state 3 after innovation/entry. The primary purpose of this graph is to display the evolution of market structure, and hence it does not necessarily convey the sense of an incumbent-entrant gap. My focus and interpretation of the historical patterns would be better summarized by the following observation. There existed 11 incumbents and zero

TABLE 2
DESCRIPTIVE STATISTICS: AGGREGATE SALES OF 5.25- AND 3.5-INCH HDDS

Variable	Observations	Mean	Standard Deviation	Minimum	Maximum
Shipment (million units)	405	1,403	3,339	.1	28,332
Average price (1998 constant US$)	405	882	892	72	7,487
Average quality (information storage capacity in megabytes)	405	2,122	5,922	15	60,000
Diameter 3.5 inches	405	.49	.50	0	1

SOURCE.—*DISK/TREND Reports*.
NOTE.—Unit of observation is product category (diameter-quality pairs), aggregated across firms and models, by year and segment.

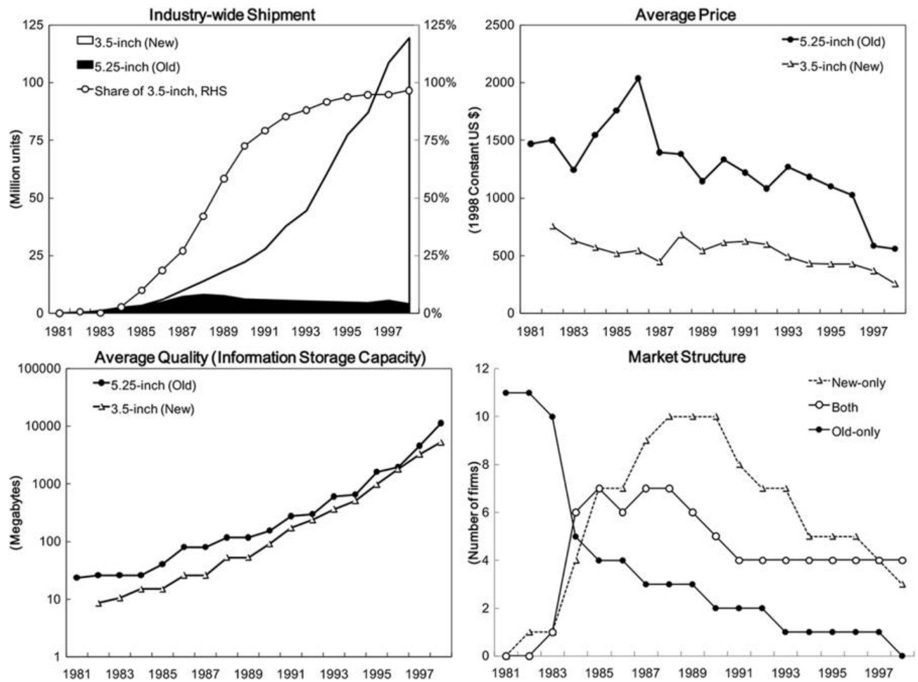

FIG. 3.—Quantity, price, quality, and industry composition. Both 5.25- and 3.5-inch HDDs serve the same market, namely, desktop PCs. Quality is measured by average capacity per unit for each generation. "Old-only" and "new-only" firms produce 5.25- and 3.5-inch HDDs, respectively. "Both" represents incumbents 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).

(actual) entrants in 1981. From a purely technological viewpoint, manufacturing of the new HDDs should have been easier for incumbents than for entrants, because certain engineering commonalities existed across multiple generations of HDDs. Nevertheless, by 1990, there were only eight innovators among incumbents, whereas at least 12 entrants started shipping new HDDs.[19] If the new HDD market could accommodate more than 11 active firms, why did incumbents not innovate as aggressively as entrants?

Section IV.A (demand estimation) explains details of the sales data. Online appendix A.1 features descriptive analysis of firm and product heterogeneity to motivate modeling choices.

III. Model

A. *Timing*

Time is discrete with finite horizon $t = 0, 1, 2, \ldots, T$. This is an important modeling choice that permits the solution of a dynamic game without ignoring the presence of fundamental nonstationarity in the data, which is a defining characteristic of an innovative industry such as high-tech manufacturing. To accommodate nonstationarity of demand and cost (increasing and decreasing with time, respectively), I allow values and policies to depend on time. A fully infinite-horizon setup will not function properly in my nonstationary data context, because demand will approach infinity and cost will approach zero, so that both the equilibrium profits and values will explode. To avoid this problem, I make a simplifying assumption that demand, cost, and market structure will stay constant after the sample period. I then associate a terminal, continuation value to each state in 1998, the terminal year of the data (see Sec. III.D for further details and empirical considerations).

A finite number of firms are indexed by i. In any year t, each firm is in one of the four technological states, $s_{it} \in$ {old only, both, new only, potential entrant}, and the industry state is their aggregation, $s_t \equiv \{s_{it}\}_i = (N_t^{\text{old}}, N_t^{\text{both}}, N_t^{\text{new}}, N_t^{\text{pe}})$, where N_t^{type} is the number of firms that belong to that type in year t, and $s_{-it} \equiv \{s_{jt}\}_{j \neq i}$. The first two types, old only and both, represent the production status of incumbent firms before and after the introduction of the new process to make the new-generation products, because pre-innovation incumbents make only old-generation products, whereas postinnovation incumbents can (but do not have to) make both old and new-generation products. The latter two types, new only and potential entrants, represent the production status of actual and poten-

[19] These numbers include firms that exited within a few years and therefore do not necessarily match the concurrent number of active firms in the graph.

tial entrants, because potential entrants become actual entrants by introducing the new-generation process.[20]

The transition of s_{it} is as follows. The game starts in year 0 with $N_0^{\text{old}} > 0$ pre-innovation incumbents, zero postinnovation incumbents ($N_0^{\text{both}} = 0$), zero actual entrants ($N_0^{\text{new}} = 0$), and $N_0^{\text{pe}} > 0$ potential entrants. I set $N_t^{\text{pe}} = 4$ for all t in the subsequent empirical implementation (see the sensitivity analysis in table 5 below for alternative specifications). In each year, a pre-innovation incumbent may either exit the industry forever, continue producing old products only, or innovate by paying a sunk cost κ^{inc} to start producing both old and new products from the next year. A postinnovation incumbent chooses to either exit or stay in the industry. A potential entrant chooses to either pay a sunk cost κ^{ent} to become an actual entrant and produce new products from the next year or quit the prospect of entry once and for all. An actual entrant chooses to exit or stay.

The timing of the game is as follows. Each year t starts with period competition among the current population of firms, s_t, from which each firm earns period profit $\pi_t(s_{it}, s_{-it})$ given the industrywide demand and cost conditions (embodied by a time subscript in $\pi_t(\cdot)$, to be specified in Sec. IV).[21] All of these industrywide features are common knowledge.

- After the period competition, N_t^{old} pre-innovation incumbents draw iid private cost shocks $\varepsilon_{it}^{\text{old}} = (\varepsilon_{it}^0, \varepsilon_{it}^1, \varepsilon_{it}^2)$ and simultaneously take actions $a_{it}^{\text{old}} \in \{\text{exit, stay, innovate}\}$.
- Having observed these actions, N_t^{both} postinnovation incumbents (excluding those incumbents that have just decided to innovate in the above) draw iid private cost shocks $\varepsilon_{it}^{\text{both}} = (\varepsilon_{it}^0, \varepsilon_{it}^1)$ and simultaneously take actions $a_{it}^{\text{both}} \in \{\text{exit, stay}\}$.
- Having observed these actions, N_t^{new} actual entrants draw iid private cost shocks $\varepsilon_{it}^{\text{new}} = (\varepsilon_{it}^0, \varepsilon_{it}^1)$ and simultaneously take actions $a_{it}^{\text{new}} \in \{\text{exit, stay}\}$.
- Having observed these actions, N_t^{pe} potential entrants draw iid private cost shocks $\varepsilon_{it}^{\text{pe}} = (\varepsilon_{it}^0, \varepsilon_{it}^1)$ and simultaneously take actions $a_{it}^{\text{pe}} \in \{\text{quit, enter}\}$.
- On the basis of these actions of firms, market structure transits from s_t to s_{t+1}. The demand and cost conditions evolve exogenously.

The order of move in the above represents another important assumption of the model to facilitate the computation of its solution as well as

[20] One can imagine entrants with the old technology as well, but I choose not to model this possibility because no such cases are recorded in the data.

[21] Potential entrants in year t do not participate in period competition in year t, and hence they are irrelevant to active firms' period profits. I use s_t (which includes N_t^{pe}) to characterize period competition and profits only for the sake of notational simplicity.

estimation. Because different types of firms move sequentially, each firm is effectively solving a single-agent problem at its turn (see Sec. III.D for further details). An empirical motivation for this specific order of move is the relative sizes and visibility of different types of firms. By definition, incumbents already exist in the market as active manufacturers and are generally larger corporate entities than entrants, which makes incumbents more conspicuous. Likewise, actual entrants are shipping HDDs whereas potential entrants are merely assessing the prospect of entry. Thus I believe that the definitional asymmetry between incumbents and entrants makes such an order-of-move assumption a natural specification for the baseline model (see cols. 2 and 3 in table 4 below for the robustness check on this assumption).

Private cost shocks reflect each firm's informational, managerial, and organizational conditions of transient nature. I focus on anonymous, type-symmetric pure strategy, which maps these cost draws to a discrete choice, in the spirit of a static entry game with private information à la Seim (2006). To facilitate both the solution and the estimation of the model, I assume that $\varepsilon_{it}(a_{it})$ is iid extreme value.

Besides the variable costs of production, active firms have to pay the fixed cost of operation, ϕ, which reflects the need for continual investment in technologies and production facilities to keep up with the industrywide trend of quality improvement: Kryder's law.[22] I set scrap values to zero because of this fast rate of obsolescence.

B. *Period Profit*

Each year, the demand and production cost conditions (D_t, C_t), the firm's own technological status (s_{it}), and the other firms' technological status (s_{-it}) completely determine the firm's period profit,

$$\pi_{it} = \pi_t^{\text{type}}(s_t) = \pi(s_{it}, s_{-it}; D_t, C_t). \tag{1}$$

The demand system D_t provides a mapping between the aggregate prices and quantities of old and new products, the substitution pattern of which will determine the relevance of cannibalization for innovating incumbents. The cost function reflects the relationship between each firm's outputs and total production costs. Section IV specifies D_t and C_t.

The HDDs are high-tech commodities with limited scope for differentiation besides product category, and hence I assume Cournot competition among (potentially) multiproduct firms, which may produce either

[22] Kryder's law says that the areal density (and hence information storage capacity) of an HDD doubles every 13 months, which is faster than Moore's law in the semiconductor industry (i.e., the circuit density of chips doubles every 18 months). The analysis of Kryder's law is outside the scope of this paper. See Lerner (1997) for the related empirical analysis.

one or both of the old- and new-generation HDDs, and I focus on anonymous, type-symmetric Nash equilibrium in the spot market.[23] Thus the market structure (summarized by the industry state s_t), along with D_t and C_t, completely determines each firm's equilibrium profit from period competition. This formulation allows us to handle the dynamic oligopoly game of innovation and entry/exit in a parsimonious state space, despite a considerably higher number of firms in the data (up to two dozen) than in typical applications of a dynamic game (between two and four).

C. *Dynamic Optimization*

When their turns to move arrive, firms make their dynamic discrete choices of entry, exit, and innovation to maximize their expected values. They discount their future stream of profits by a factor $\beta \in (0, 1)$, with rational expectations regarding the endogenous evolution of market structure and perfect foresight regarding the exogenous evolution of demand and production costs.[24]

These assumptions are strong but useful for three reasons. First, the existing explanations for the innovator's dilemma, such as in Christensen (1993, 1997), tend to rely on the assumptions of asymmetric irrationality. Researchers have argued that the managers at incumbent firms suffer from cognitive biases and other informational problems, but they typically assume more rational beliefs for entrants, thereby trying to explain the seemingly suboptimal investment behaviors of incumbents by ad hoc assumptions of irrational beliefs. By contrast, this paper aims to offer rational and less ad hoc explanations, and hence these assumptions are consistent with the purpose of this research. Second, the evolution of demand and costs in the HDD industry has historically followed regular patterns. Demand grew steadily with the expansion of the computer market; the manufacturing costs decreased steadily as a result of Kryder's law (see the end of Sec. III.A). Third, these assumptions enhance tractability and reduce the computational burden in the estimation of the model. Thus, these assumptions are integral parts of the research design.[25]

[23] Another motivation for the Cournot competition is that production facilities take time to build, up to a year. Hence we can invoke Kreps and Scheinkman's (1983) argument that capacity building followed by pricing leads to Cournot outcomes.

[24] I assume that firms know the entire history of $\{(D_t, C_t)\}_t$ from the beginning.

[25] These are the motivations for my assumptions and do not imply their innocuousness. The degree of uncertainty about demand and costs can have important implications for the amount and timing of investment, as Dixit and Pindyck (1994) pointed out. See Collard-Wexler (2013) for evidence on this point based on the estimation of a structural dynamic oligopoly game. In the current context, however, the existence of uncertainty alone would not influence the key empirical finding on the incumbent-entrant gap. Uncertainty leads to inaction and delay of innovation, but both incumbents and entrants are operating in the same environ-

The model incorporates preemptive motives as a part of dynamic strategic interactions, in the sense that the firm's own entry or innovation makes the future market more competitive and therefore makes the other firms' subsequent entry or innovation less profitable and less likely. Because no additional channels exist through which the first-mover advantage manifests itself, the model focuses on the simplest notion of preemptive motives based on the pure market structure effect.[26]

The dynamic programming problems of active firms are characterized by the following Bellman equations:

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

$$\begin{aligned} V_t^{\text{both}}(s_t, \varepsilon_{it}) = {} & \pi_t^{\text{both}}(s_t) \\ & + \max\{\varepsilon_{it}^0, -\phi + \beta E[V_{t+1}^{\text{both}}(s_{t+1}, \varepsilon_{it+1})|s_t, \varepsilon_{it}] + \varepsilon_{it}^1\}, \end{aligned} \tag{3}$$

and

$$\begin{aligned} V_t^{\text{new}}(s_t, \varepsilon_{it}) = {} & \pi_t^{\text{new}}(s_t) \\ & + \max\{\varepsilon_{it}^0, -\phi + \beta E[V_{t+1}^{\text{new}}(s_{t+1}, \varepsilon_{it+1})|s_t, \varepsilon_{it}] + \varepsilon_{it}^1\}, \end{aligned} \tag{4}$$

subject to the perceived law of motion governing s_t. The expectations are over the other firms' choices and hence over the realizations of their private cost shocks. For a potential entrant, the problem is simply

$$\max\{\varepsilon_{it}^0, \beta E[V_{t+1}^{\text{new}}(s_{t+1}, \varepsilon_{it+1})|s_t, \varepsilon_{it}] - \kappa^{\text{ent}} + \varepsilon_{it}^1\}. \tag{5}$$

Besides the components of period profit functions, the key parameters of this dynamic discrete game are the sunk cost of innovation for incumbents, κ^{inc}, that of entry/innovation for potential entrants, κ^{ent}, and the

ment and hence share this uncertainty in common. Unless incumbents faced systematically more uncertainty than entrants (which is the opposite of the conventional characterization), uncertainty cannot explain the innovation gap. Moreover, any systematic difference between incumbents and entrants (including heterogeneous beliefs and cognitive biases about the net benefit from innovation), if it exists in the data, should manifest itself as differential sunk costs of innovation in my estimation results. These cost estimates represent heterogeneous R&D efficiencies and absorb any innovation gap that is not explained by the model. To use a production function analogy, the role of κ is just like the Solow residual, which picks up total factor productivity as well as everything else outside the model. See Secs. VI.C and VII.A for further discussions.

[26] Other channels may also exist in reality, such as time to build production capability and recognition among buyers, and hence I view my subsequent empirical results as the lower bound of preemptive motives.

fixed cost of operation, ϕ.[27] Thus the model incorporates the incumbent-entrant heterogeneity in the efficiency of innovation and allows incumbents to possess either advantages or disadvantages relative to potential entrants.

D. *Equilibrium*

I solve this finite-horizon, sequential-move dynamic discrete game with private information for a PBE in type-symmetric pure strategies. Three features of the model are important to ensure computational feasibility and avoid multiple equilibria. First, because private information is merely in the form of iid cost shocks associated with each firm's discrete alternatives, $\varepsilon(a_{it})$, and not in the form of persistent heterogeneity, the firm's belief over off-path realizations of $\varepsilon(a_{-it})$ does not affect its payoff.[28] That is, the firm's payoff is affected by its rivals' cost shocks only through their actual choices, and not by the specific realizations of $\varepsilon(a_{-it})$, so firms hold perfect information on the payoff-relevant part of past history. Second, different types of firms move sequentially after observing the entry/exit/innovation choices of earlier movers. At its turn to move, the firm (or the same type of firms with symmetric strategies) is effectively solving a single-agent problem based on its expectation over the subsequent evolution of market structure. Third, these two features and the finite-horizon formulation allow us to solve the model by backward induction.

I assume that the terminal values associated with a firm's states, $s_{iT} \in$ {old, both, new}, are[29]

$$\begin{aligned}&\left(V_T^{\text{old}}, V_T^{\text{both}}, V_T^{\text{new}}\right)\\ &= \left(\sum_{\tau=T}^{\infty}\beta^{\tau}\pi_T^{\text{old}}(s_T), \sum_{\tau=T}^{\infty}\beta^{\tau}\pi_T^{\text{both}}(s_T), \sum_{\tau=T}^{\infty}\beta^{\tau}\pi_T^{\text{new}}(s_T)\right).\end{aligned} \tag{6}$$

[27] I normalize the scrap value upon exit to zero and omit it from the model because *DISK/TREND Report* rarely indicates any profitable sales of facilities or equipment when firms exit the market, which seems consistent with the industry's fast pace of obsolescence.

[28] Because PBE and sequential equilibrium (SE) differ only in terms of restrictions on off-path beliefs, we may alternatively use SE as a solution concept for the same results.

[29] I am reconciling the finite-horizon model with the reality in which the world did not actually end in 1998, by assuming that the state stops evolving after year T. These terminal values would reflect an analyst's assumptions on the postsample periods. The model's solution and parameter estimates will depend on these assumptions. However, the estimated model's qualitative implications will not depend much on them. Because the demand for 5.25-inch HDDs had all but disappeared by 1998 and $\pi_t(s_T)$'s are pinned down by the static parameters as well as the Cournot competition assumption, V_T's play only a limited role as a scaling parameter of the game's payoffs. 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.

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

$$\max\{\varepsilon^0_{i,T-1}, -\phi + \beta E[V^{\text{old}}_T(s_T)|s_{T-1}] + \varepsilon^1_{i,T-1},$$
$$-\phi + \beta E[V^{\text{both}}_T(s_T)|s_{T-1}] - \kappa^{\text{inc}} + \varepsilon^2_{i,T-1}\}.$$

I follow Rust (1987) to exploit the property of the logit errors, $\varepsilon_{it}(a_{it})$, and their conditional independence over time, to obtain a closed-form expression for the expected value before observing $\varepsilon_{it}(a_{it})$,

$$\begin{aligned} &E_{\varepsilon_{i,T-1}}[V^{\text{old}}_{T-1}(s_{T-1}, \varepsilon_{i,T-1})|s_{T-1}] \\ &\quad = \pi^{\text{old}}_{T-1}(s_{T-1}) + \gamma + \ln(\exp(0) \\ &\qquad + \exp\{-\phi + \beta E[V^{\text{old}}_T(s_T)|s_{T-1}]\} \\ &\qquad + \exp\{-\phi + \beta E[V^{\text{both}}_T(s_T)|s_{T-1}] - \kappa^{\text{inc}}\}), \end{aligned}$$

where γ is the Euler constant. Similar expressions hold for the other two types.[30] 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 maximum likelihood estimation (MLE).

IV. 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 s_t. Third, I embed these period profits into the dynamic discrete game of innovation and entry/exit, which I solve to estimate the sunk costs of innovation, entry, and continued operation.

[30] The ex ante values for a postinnovation incumbent and an actual entrant are as follows:

$$E_{\varepsilon_{i,T-1}}[V^{\text{both}}_{T-1}(s_{T-1}, \varepsilon_{i,T-1})|s_{T-1}]$$
$$= \pi^{\text{both}}_{T-1}(s_{T-1}) + \gamma + \ln(\exp(0) + \exp\{-\phi + \beta E[V^{\text{both}}_T(s_T)|s_{T-1}]\}),$$
$$E_{\varepsilon_{i,T-1}}[V^{\text{new}}_{T-1}(s_{T-1}, \varepsilon_{i,T-1})|s_{T-1}]$$
$$= \pi^{\text{new}}_{T-1}(s_{T-1}) + \gamma + \ln(\exp(0) + \exp\{-\phi + \beta E[V^{\text{new}}_T(s_T)|s_{T-1}]\}).$$

A. Demand

I capture the substitution pattern across generations of HDDs using the multinomial logit model of differentiated products. Although the use of a discrete-choice model for demand analysis is a common practice, note that this paper's application departs from the standard notational convention of denoting firm or brand by a j subscript, because my model focuses on product differentiation across categories but not firms, because of the homogeneous "commodity" nature of HDDs. The dynamic oligopoly game framework in the previous section highlights HDDs' differentiation across generations and assumes homogeneity within each generation. The empirical demand analysis incorporates more details to exploit additional variations in the data, in which the unit of observation is the combination of generation, quality, buyer category, geographical regions, and year t. 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[31]

$$u_{kj} = \alpha_0 + \alpha_1 p_j + \alpha_2 I(g_j = \text{new}) + \alpha_3 x_j + \xi_j + \epsilon_{kj}, \tag{7}$$

with a j subscript denoting product category (not firm or brand), where p_j is the price, ξ_j is the unobserved characteristics (most importantly, "design popularity" among buyers, as well as other unobserved attributes such as "reliability"), and ϵ_{kj} is the idiosyncratic taste shock that is assumed iid extreme value (over buyers and generation-quality bins). The outside goods offer the normalized utility $u_{k0} \equiv 0$, which represent *removable* HDDs (as opposed to *fixed* HDDs) and other storage devices (e.g., tape recorders, optical disk drives, and flash memory).

Let $\bar{u}_j \equiv \alpha_0 + \alpha_1 p_j + \alpha_2 I(g_j = \text{new}) + \alpha_3 x_j + \xi_j$ represent the mean utility from a category j HDD whose market share is $ms_j = \exp(\bar{u}_j)/\sum_l \exp(\bar{u}_l)$. The shipment quantity is $Q_j = ms_j M$, where M is the size of the HDD market including the outside goods. Practically, M reflects all desktop PCs to be manufactured globally in a given year. Berry's (1994) inversion provides the linear relationship

[31] I suppress the time subscript t for simplicity of notation. The demand side is static in the sense that buyers make new purchasing decisions every year. Multiyear contracting is not common, and hundreds of buyers (e.g., computer makers) are present during the sample period. I do not model HDDs as durable goods because of fast obsolescence due to Kryder's law and also because the dynamics of repurchasing cycles in the PC market is driven primarily by operating systems (e.g., Windows 95 and 98) or central processing unit chips (e.g., Intel's Pentium III), which I assume evolve exogenously to the HDD market. See online app. A.1

$$\ln\left(\frac{ms_j}{ms_0}\right) = \alpha_1 p_j + \alpha_2 I\left(g_j = \text{new}\right) + \alpha_3 x_j + \xi_j, \tag{8}$$

where ms_0 is the market share of outside goods (removable HDDs). I estimate the taste parameters (α_1, α_2, α_3) by instrumental variable (IV) regressions of this linear equation.

Sources of identification.—The demand parameters are identified by the time-series and cross-sectional variations in the data.[32] 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 reflected in *DISK/TREND Reports*. Second, data are recorded by buyer category, PC makers, and distributors/end users. Third, data are recorded by geographical category, US and non-US.

In the IV estimation, I use the following variables as instruments for p_j: (1) the prices in the other region and user category and (2) the number of product "models" (not firms). The first IV is used by Hausman (1996) and 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 is used by Bresnahan (1981) and Berry, Levinsohn, and Pakes (1995) and exploits the proximity of rival products (in product space), that is, the negative correlation between markup and the number of "models" in oligopolies. The identifying assumption is that taste shocks (i.e., ξ_{jt}) in any given period are not correlated with the number of models in a particular product category j, which are outside my model.[33]

These two 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 alternative, time-series orthogonality conditions in the style of Aguirregabiria and Ho (2012) and Sweeting (2013) and obtained the price coefficient estimates of approximately −3.20, a range statistically indistinguishable at the 5 percent level from my preferred estimate of −3.28 based on the other three IVs (see Sec. V.A, col. 4 of table 3 below).

[32] See Berry and Haile (2009) for nonparametric identification of static discrete-choice demand models, using the types of instruments I use in the following.

[33] 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.

This third approach 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 (the estimate for ρ is .41), 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 that firms at t do not know the unpredictable parts ν_{jt+1} when they make dynamic decisions.[34]

B. *Period Competition and Marginal Costs*

Multiproduct (i.e., old and new goods) Cournot competition characterizes the spot market competition.[35] 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

$$\pi_i = \sum_{g\in A_i} \pi_{ig} = \sum_{g\in A_i} \left(p_g - mc_g\right) q_{ig} \tag{9}$$

with respect to shipping quantity q_{ig} for all $g \in A_i$, 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

$$p_g + \frac{\partial p_g}{\partial Q_g} q_{ig} + \frac{\partial p_h}{\partial Q_g} q_{ih} = mc_g, \tag{10}$$

with $g, h \in \{\text{old, new}\}$, $g \neq h$, 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.

For each year, we can infer the marginal costs of production, mc_{old} and mc_{new}, from equation (10). Because the unit of 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.

[34] 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).

[35] Besides the data constraint described in Sec. II.B, three additional considerations motivate 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 à la Kreps and Scheinkman (1983). Third, accounting records indicate that despite fierce competition with undifferentiated goods, the HDD makers seemed to enjoy nonzero (albeit razor-thin) profit margins on average.

C. Costs of Innovation, Entry, and Continued Operation

These static demand and cost estimates from the previous two steps imply specific period profit for each type of firms, in each year, under each market structure. In the third and final step of estimation, I embed these variable profits into the dynamic discrete game model and solve it for a PBE by backward induction (see Sec. III.D for details). The goal of this step is to obtain estimates for the three dynamic parameters, $(\phi, \kappa^{\text{inc}}, \kappa^{\text{ent}})$, by maximum likelihood. Given a vector of candidate parameter values, I can solve the dynamic game. That is, each possible combination of $(\phi, \kappa^{\text{inc}}, \kappa^{\text{ent}})$ implies a specific expected value for each firm type, in each state-year, as well as the optimal choice probabilities of entry/exit and innovation. The ML estimate is the vector that maximizes the likelihood of observing the actual choice probabilities in the data.

By contrast, I do not intend to estimate the discount factor, β, because its identification is known to be impractical (cf. Rust 1987). Likewise, although an additional parameter, the rate of change in sunk costs, δ, is desirable for a better fit of entry timing patterns, δ turns out to be difficult to estimate; so instead I will assume that δ equals some constant and subsequently conduct sensitivity analysis (Sec. V.C).

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

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

where $pr_t^{\text{old}}(\cdot)$ is the time-specific probability that an old-only firm takes a particular action d_{it}:[36]

$$\begin{aligned} pr_t^{\text{old}}(d_{it} = \text{exit}) &= \exp(0)/B, \\ pr_t^{\text{old}}(d_{it} = \text{stay}) &= \exp[-\phi + \beta E_\varepsilon V_{t+1}^{\text{old}}(s_{t+1})]/B, \\ pr_t^{\text{old}}(d_{it} = \text{adopt}) &= \exp[-\phi + \beta E_\varepsilon V_{t+1}^{\text{both}}(s_{t+1}) - \delta^t \kappa^{\text{inc}}]/B, \end{aligned}$$

where

$$\begin{aligned} B \equiv{} & \exp(0) + \exp[-\phi + \beta E_\varepsilon V_{t+1}^{\text{old}}(s_{t+1})] \\ & + \exp[-\phi + \beta E_\varepsilon V_{t+1}^{\text{both}}(s_{t+1}) - \delta^t \kappa^{\text{inc}}]. \end{aligned}$$

The contributions of the other three types of firms take similar forms (see online app. A.2).

Year t has $N_t \equiv (N_t^{\text{old}}, N_t^{\text{both}}, N_t^{\text{new}}, N_t^{\text{pe}})$ active firms in each state, of which $X_t \equiv (X_t^{\text{old}}, X_t^{\text{both}}, X_t^{\text{new}})$ exit and $E_t \equiv (E_t^{\text{old}}, E_t^{\text{pe}})$ innovate. Denote the joint likelihood for year t of observing data (N_t, X_t, E_t) by $P(N_t, X_t,$

[36] See online app. A.5.1 for methodological considerations behind these nonstationary policy functions.

E_t). Then the overall joint likelihood for $t = 0, 1, 2, \ldots, T - 1$ is $P(N, X, E) = \prod_{t=0}^{T-1} P(N_t, X_t, E_t)$. Thus the ML estimators for the mean fixed cost of operation ϕ and the base sunk costs of innovation/entry κ^{inc} and κ^{ent} are

$$\arg\max_{\phi,\kappa^{\text{inc}},\kappa^{\text{ent}}} \ln[P(N, X, E)]. \tag{11}$$

Sources of identification.—I obtain the static demand and cost estimates (and hence the implied period profits, $\pi_t^{\text{type}}(s_t)$) from the HDD sales data and completely outside the dynamic estimation framework. As such, these static estimates, together with the observed entry/exit/innovation choices in the panel data of HDD manufacturers, constitute the key inputs for identifying the dynamic parameters. For example, a large fixed cost ϕ will decrease the predicted value of $pr_t^{\text{old}}(d_{it} = \text{stay})$, $pr_t^{\text{old}}(d_{it} = \text{adopt})$, and B and hence increase the predicted optimal choice probability of exit, $pr_t^{\text{old}}(d_{it} = \text{exit})$. Correspondingly, if a high fraction of active firms actually choose to exit in the panel data, such a data pattern (i.e., high $X_t^{\cdot}/N_t^{\cdot}$) will lead to a large estimate of $\hat{\phi}$. Likewise, large sunk costs of innovation, κ^{inc} and κ^{ent}, will decrease $pr_t^{\text{old}}(d_{it} = \text{adopt})$ and $pr_t^{\text{pe}}(d_{it} = \text{enter})$, respectively, so the observed fractions of innovating incumbents and potential entrants in the data (i.e., $E_t^{\text{old}}/N_t^{\text{old}}$ and $E_t^{\text{pe}}/N_t^{\text{pe}}$) will differentially pin down the MLEs for $\hat{\kappa}^{\text{inc}}$ and $\hat{\kappa}^{\text{ent}}$.[37]

V. Results

A. *Demand*

Table 3 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 (US and non-US) 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 data set 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 statistically significant than the ordinary least squares (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 mar-

[37] See Aguirregabiria and Suzuki (2014) for a formal identification discussion on dynamic entry models.

TABLE 3
LOGIT DEMAND ESTIMATES FOR 5.25- AND 3.5-INCH HDDS

	BROAD MARKET		NARROW MARKET	
	OLS (1)	IV (2)	OLS (3)	IV (4)
Price ($000)	−1.66***	−2.99***	−.93**	−3.28***
	(.45)	(.55)	(.46)	(.63)
Diameter 3.5 inches	.84*	.75	1.75***	.91**
	(.46)	(.45)	(.31)	(.38)
Log capacity (megabytes)	.18	.87***	.04	1.20***
	(.33)	(.27)	(.26)	(.31)
Year dummies	Yes	Yes	Yes	Yes
Region/user dummies	. . .	. . .	Yes	Yes
Adjusted R^2	.43	.33	.50	.28
Observations	176	176	405	405
Partial R^2 for price	. . .	.32	. . .	.16
p-value	. . .	.00	. . .	.00

NOTE.—Standard errors, in parentheses, are clustered by capacity-diameter (cols. 1 and 2) and capacity-diameter-region-user (cols. 3 and 4).
* Significant at the 10 percent level.
** Significant at the 5 percent level.
*** Significant at the 1 percent level.

ket definition, as my baseline result for the subsequent analyses, because 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_j in eqq. [7] and [8]). I use equation (8) to recover $\hat{\xi}_{jt}$ as residuals. Figure 4 (left panel) shows the evolution of $\hat{\xi}_{jt}$ for both old and new HDDs, the positions of which switched in 1992 and suggest that the 3.5-inch replaced the 5.25-inch as the mainstream HDD type.

B. Marginal Costs

From the demand estimates and firms' first-order conditions, I infer marginal costs of production (fig. 4, right panel).[38] The continual drop in the marginal costs reflects two tendencies in the industry. First, HDDs required increasingly fewer parts because of design improvements, reflecting advances

[38] My model assumes that the marginal costs, mc_{old} and mc_{new}, are common to all types of firms, which precludes economies of scope due to the joint production of both the old and new HDDs. One might suspect the existence of some economies of scope due to shared inputs or processes. Unfortunately, the original data source does not report the firm-level market shares by product category for confidentiality reasons, and hence I cannot investigate this possibility by estimating marginal costs by firm type.

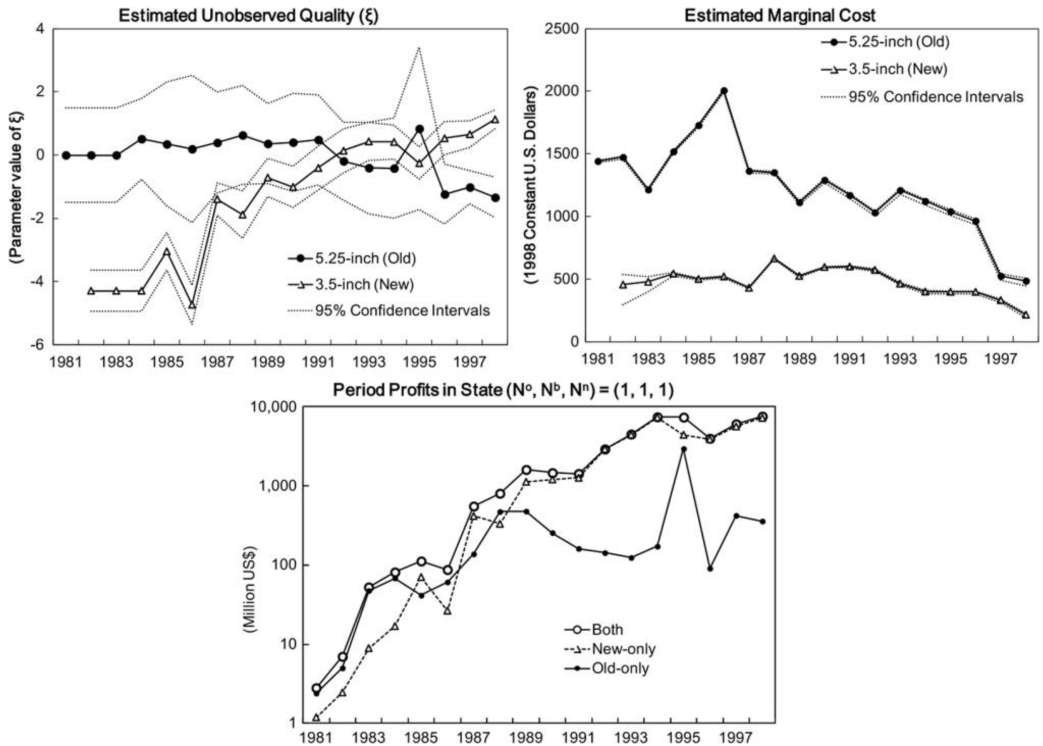

FIG. 4.—Estimated unobserved quality, marginal costs, and period profits. Results based on the IV estimates of logit demand system with the narrow market definition (col. 4 of table 3).

in engineering. Second, offshore production in Singapore and other Southeast Asian locations became prevalent, reducing primarily the cost of hiring engineers. Together these developments represent important channels of process innovation.[39]

Figure 4 (bottom panel) shows the estimated period profits, for a constant market structure $(N_t^{\text{old}}, N_t^{\text{both}}, N_t^{\text{new}}) = (1, 1, 1)$, to illustrate how the incentives to innovate have been changing over time. The most salient feature is the rapid growth of profits thanks to the explosion of the demand for PCs (fig. 3, top left), which is the reason I use a log scale. However, the profit for the old-only firm stopped growing in 1990 as the new HDDs became mainstream. An upward spike in 1995 is an exception due to the release of Windows 95, which triggered a temporary shortage of various components of PCs. The most important pattern concerning the incentives to innovate is that an incumbent would earn significantly higher profits from producing both the old and new HDDs than as an old-only producer ($\pi_t^{\text{both}} > \pi_t^{\text{old}}$ for all t).

C. *Fixed and Sunk Costs*

Table 4 shows the MLEs of the mean fixed cost of operation, ϕ, and the base sunk costs of innovation, κ^{inc} and κ^{ent}, under different assumptions on the order of moves. The estimates suggest that the cost of innovation is lower for incumbents than for entrants ($\hat{\kappa}^{\text{inc}} < \hat{\kappa}^{\text{ent}}$); therefore, the seeming inertia of incumbents does not stem from their innate cost disadvantage.[40] The explanation lies in other incentives, which I explore in detail with counterfactual analyses in Section VI.

My baseline model in column 1 specifies the following order of moves: old-only, both, new-only, and potential entrants (see Sec. III.D for the reasons). As a robustness check, column 2 shows the estimates based on the reverse ordering assumption, and the specification for column 3 prioritizes entry and innovation decisions. The overall magnitude of changes in parameter estimates is negligible, because the order of move is not as decisive a factor in a fully dynamic (multiperiod) model as in a more stylized (static or two-period) model. Nevertheless, qualitative differences

[39] See Igami (forthcoming) for details on the industry dynamics of offshoring.

[40] The result $\hat{\kappa}^{\text{inc}} < \hat{\kappa}^{\text{ent}}$ does not necessarily mean that incumbents are entirely free from organizational, informational, or other disadvantages. Rather, my estimates simply suggest that 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 $\hat{\kappa}^{\text{inc}}$ and $\hat{\kappa}^{\text{ent}}$ is beyond the scope of this paper because of limited data on firms' internal affairs, but it will be an interesting topic for future research in its own right.

TABLE 4
MAXIMUM LIKELIHOOD ESTIMATES OF THE DYNAMIC PARAMETERS

	ASSUMED ORDER OF MOVES		
	Old-Both-New-PE (1)	PE-New-Both-Old (2)	PE-Old-Both-New (3)
Fixed cost of operation (ϕ)	.1474	.1472	.1451
	[−.02, .33]	[−.02, .33]	[−.03, .33]
Incumbents' sunk cost (κ^{inc})	1.2439	1.2370	1.2483
	[.51, 2.11]	[.50, 2.10]	[.51, 2.11]
Entrants' sunk cost (κ^{ent})	2.2538	2.2724	2.2911
	[1.74, 2.85]	[1.76, 2.87]	[1.78, 2.89]
Log likelihood	−112.80	−112.97	−113.46

NOTE.—Figures are in billions of dollars. The 90 percent confidence intervals are in brackets. These intervals are based on likelihood ratio tests (i.e., by comparing $LL(\hat{\theta}) - LL(\tilde{\theta})$ with the corresponding critical values from the χ^2 distribution, where $LL(\theta)$ is the joint log likelihood, ln $P(N, X, E)$, evaluated at a vector of parameter values $\theta \equiv (\phi, \kappa^{\text{inc}}, \kappa^{\text{ent}})$; $\hat{\theta}$ and $\tilde{\theta}$ represent the MLE and its perturbed counterpart, respectively), and represent the lower bounds because they do not incorporate the error from the estimation of static parameters. Standard errors based on Wald-type tests are either implausibly small or large because the likelihood function exhibits a step function–like shape due to the game-theoretic nature of the model.

appear consistent with the economic intuition about early-mover advantages.[41]

The cost estimates on the order of billions of dollars might appear implausibly high at first glance, but they are comparable to the annual R&D budget at specialized HDD manufacturers such as Western Digital and Seagate Technology (between \$0.6 billion and \$1.6 billion). Their capital expenditures (i.e., investments in plants and equipment) have also fluctuated within the same range.[42]

Table 5 reports the results of sensitivity analysis with respect to the discount factor, β, the rate of change of innovation costs, δ, and the number of potential entrants, N^{pe}. I set $\beta = .8$, $\delta = 1.1$, and $N_t^{\text{pe}} = 4$ for all t for my baseline estimates in table 4. All of the cost estimates $(\hat{\phi}, \hat{\kappa}^{\text{inc}}, \hat{\kappa}^{\text{ent}})$ increase with β, because a higher discount factor means higher expected

[41] For example, incumbents are more "handicapped" in model 2 than in model 1, which implies a lower sunk cost estimate $\hat{\kappa}_2^{\text{ent}} < \hat{\kappa}_1^{\text{ent}}$, because their actual innovation behavior in the data does not vary by my modeling assumptions and model 2 needs a lower cost to rationalize such a data pattern. By contrast, potential entrants are less "handicapped" in models 2 and 3 than in model 1, and hence these models generate higher cost estimates $(\hat{\kappa}_2^{\text{ent}} > \hat{\kappa}_1^{\text{ent}}$ and $\hat{\kappa}_3^{\text{ent}} > \hat{\kappa}_1^{\text{ent}})$ to rationalize the observed entry/innovation pattern.

[42] Two caveats are in order regarding this ballpark comparison of economic and accounting costs. First, not all of these accounting items relate specifically to the development of new HDDs. Second, my estimates incorporate any economic costs associated with the decision making and execution of innovation, including cognitive, informational, and organizational costs as well as direct, tangible costs of developing product designs, key components, manufacturing equipment, and a year-long process of trials and errors to achieve reliable volume production.

TABLE 5
SENSITIVITY ANALYSIS

	$\beta = .75$	$\beta = .85$	$\beta = .90$
Fixed cost of operation (ϕ)	−.04	.35	.55
Incumbents' sunk cost (κ^{inc})	1.23	1.39	1.88
Entrants' sunk cost (κ^{ent})	2.25	2.50	3.06
Log likelihood	−110.22	−116.84	−127.21
	$\delta = 1.00$	$\delta = 1.05$	$\delta = 1.15$
Fixed cost of operation (ϕ)	.12	.14	.07
Incumbents' sunk cost (κ^{inc})	1.59	1.47	1.22
Entrants' sunk cost (κ^{ent})	4.26	3.12	1.99
Log likelihood	−124.09	−116.00	−113.60
	$N^{pe} = 1$	$N^{pe} = 2$	$N^{pe} = 3$
Fixed cost of operation (ϕ)	.14	.10	.14
Incumbents' sunk cost (κ^{inc})	1.28	1.29	1.28
Entrants' sunk cost (κ^{ent})	−.06	1.83	2.09
Log likelihood	−96.94	−103.03	−108.61

NOTE.—As a reminder, the baseline parameter values for β, δ, and N^{pe} are .8, 1.1, and 4, respectively, and the corresponding log likelihood is −112.80.

values from the future operation and innovation, and hence the model needs correspondingly higher cost estimates to rationalize the observed patterns of entry, exit, and innovation. The sunk cost estimates ($\hat{\kappa}^{inc}$, $\hat{\kappa}^{ent}$) decrease with δ, because a higher growth rate of the innovation costs leads to higher average costs of innovation over time; hence the model requires correspondingly lower levels of the base sunk costs to rationalize the innovation/entry rates in the data. The value of $\hat{\kappa}^{ent}$ increases with N^{pe} because models with higher numbers of potential entrants need higher entry costs to rationalize the number of actual entrants in the data.

Given the paper's focus on the incumbent-entrant innovation gap, we should be particularly careful about N^{pe} because it affects $\hat{\kappa}^{inc}$ and $\hat{\kappa}^{ent}$ differently, whereas β and δ influence them in the same direction and are therefore relatively innocuous to my inference that $\hat{\kappa}^{inc} < \hat{\kappa}^{ent}$. I have chosen $N^{pe} = 4$ in the baseline model for the following reasons. On the one hand, four is the minimum number of potential entrants to rationalize the data, which contain the maximum of four entries per year in the mid-1980s.[43] This aspect of the data rejects any models with $N^{pe} < 4$, even though their log likelihoods might appear more attractive. On the other hand, models with $N^{pe} > 4$ will automatically lead to higher $\hat{\kappa}^{ent}$ and favor my finding that $\hat{\kappa}^{inc} < \hat{\kappa}^{ent}$ without any justification from the data. The

[43] By contrast, the data contain a zero entry during the 1990s, and the models with $N^{pe} <$ 4 fit this part of the data well, which is why their log likelihoods are higher than that in the baseline setting.

baseline estimates under $(\beta, \delta, N^{\text{pe}}) = (.8, 1.1, 4)$ led to the log likelihood of -112.8, which is higher than those from most other configurations besides those with $N^{\text{pe}} < 4$.[44] Thus I believe that the baseline calibration represents a reasonable modeling choice for the purpose of estimating the innovator's dilemma.

Table A4 in online appendix A.3.1 presents estimates of an alternative specification of the model that relaxes the assumption that the incumbents' fixed costs of operation are the same when producing one or two generations of products (i.e., an extreme form of economies of scope). Key findings include the following: (1) the sunk cost estimates are similar to the baseline results; (2) the fixed-cost estimates indicate some diseconomies of scope, although the low precision of estimates for ϕ's precludes any definite conclusions; and (3) the counterfactual experiments based on these alternative estimates lead to quantitatively similar outcomes on the innovation gap.

D. Fit

Figure 5 suggests that the estimated model fits the data reasonably well, replicating three key features of innovation and market structure dynamics, albeit in a slightly smoother manner. First, the number of old-only firms declines precipitously during the 1980s, as some of them innovate whereas others exit, and then more slowly during the 1990s. Second, the number of both firms increases until the late 1980s and then stabilizes at around four during the 1990s. Third, the number of new-only firms peaks at a much higher level and then declines at a faster rate than innovative incumbents. The estimated model does not replicate all of the wiggles in the data but appears to provide a simple benchmark against which we can compare alternative industry dynamics under different circumstances, in Sections VI and VII.

VI. Rational Innovator's Dilemma

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 gap in innovation: cannibalization, preemption, and heterogeneous sunk costs. To measure each effect, I compare the gaps in the estimated baseline model with a counterfactual simulation in which that particular incentive mechanism is absent.

[44] The model with $\beta = .75$ also entails a higher likelihood; but a negative $\hat{\phi}$ is counterintuitive for the HDD industry, in which firms incur sizable costs of operation and continual R&D investments.

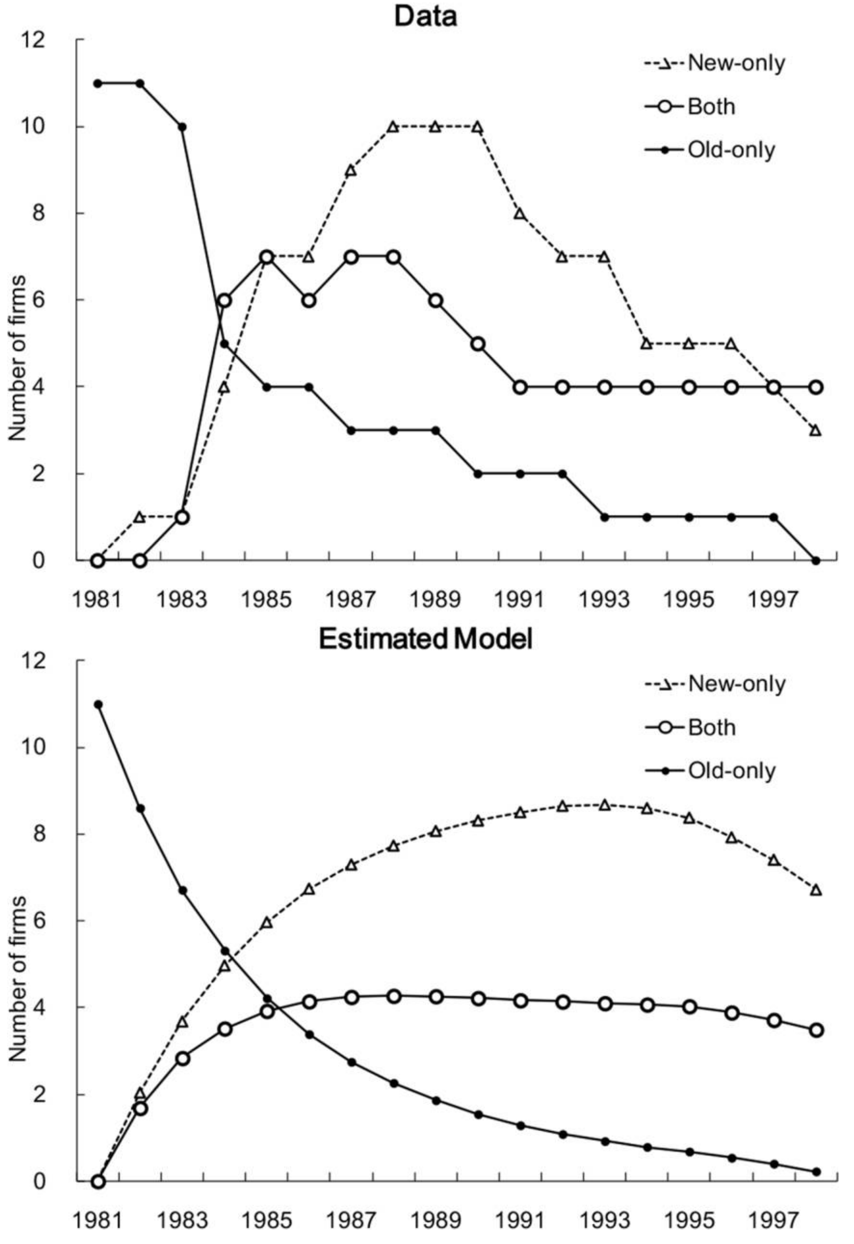

FIG. 5.—Fit of market structure dynamics. The bottom panel displays the mean numbers of firms across 10,000 simulations of the estimated model.

A. *Cannibalization*

First, I hypothetically 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.[45] 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 innovate and become both. The latter division acts like a potential entrant with staying power, which can choose to innovate (and become new-only) or wait.

Operationally, I solve the model for a new PBE in this counterfactual environment and use the equilibrium choice probabilities to run 10,000 simulations of industry history. The simulation results suggest that incumbents (their "corporate venture" divisions, to be precise) would be much more eager to innovate than in the baseline case (fig. 6, top). Free of the cannibalization concerns regarding their own old-HDD business, 8.95 incumbents start producing new HDDs in the first 10 years, compared with 6.30 in the baseline model. As a measure of the incumbent-entrant innovation gap, we can also compare the cumulative numbers of innovators among incumbents and entrants at the end of the sample period (see fig. 1, top). In the baseline model, 6.45 more entrants than incumbents have already innovated by 1998. In the no-cannibalization counterfactual, this gap would shrink by 57 percent to only 2.80 firms. Thus cannibalization can explain a significant part of the incumbent-entrant innovation gap.

B. *Preemption*

Second, I remove preemptive motives from incumbents' innovation decisions by making the rate of new entry unresponsive to incumbents' innovations, so that incumbents cannot ever hope to deter entrants by their own innovations. Specifically, I force potential entrants to ignore incumbents' innovations, that is, make entry/innovation decisions as if

$$\tilde{s}_t = \left(\tilde{N}_t^{\text{old}}, \tilde{N}_t^{\text{both}}, \tilde{N}_t^{\text{new}}, \tilde{N}_t^{\text{pe}}\right) = \left(N_t^{\text{old}} + N_t^{\text{both}}, 0, N_t^{\text{new}}, N_t^{\text{pe}}\right),$$

and focus on how incumbents would differently best respond to such hypothetical potential entrants. Put differently, potential entrants in the baseline model are rational and less likely to enter when N_t^{both} is (expected to be) higher, and incumbents rationally expect this preemptive effect

[45] See online app. A.4.1 for an alternative approach to isolating the cannibalization concern.

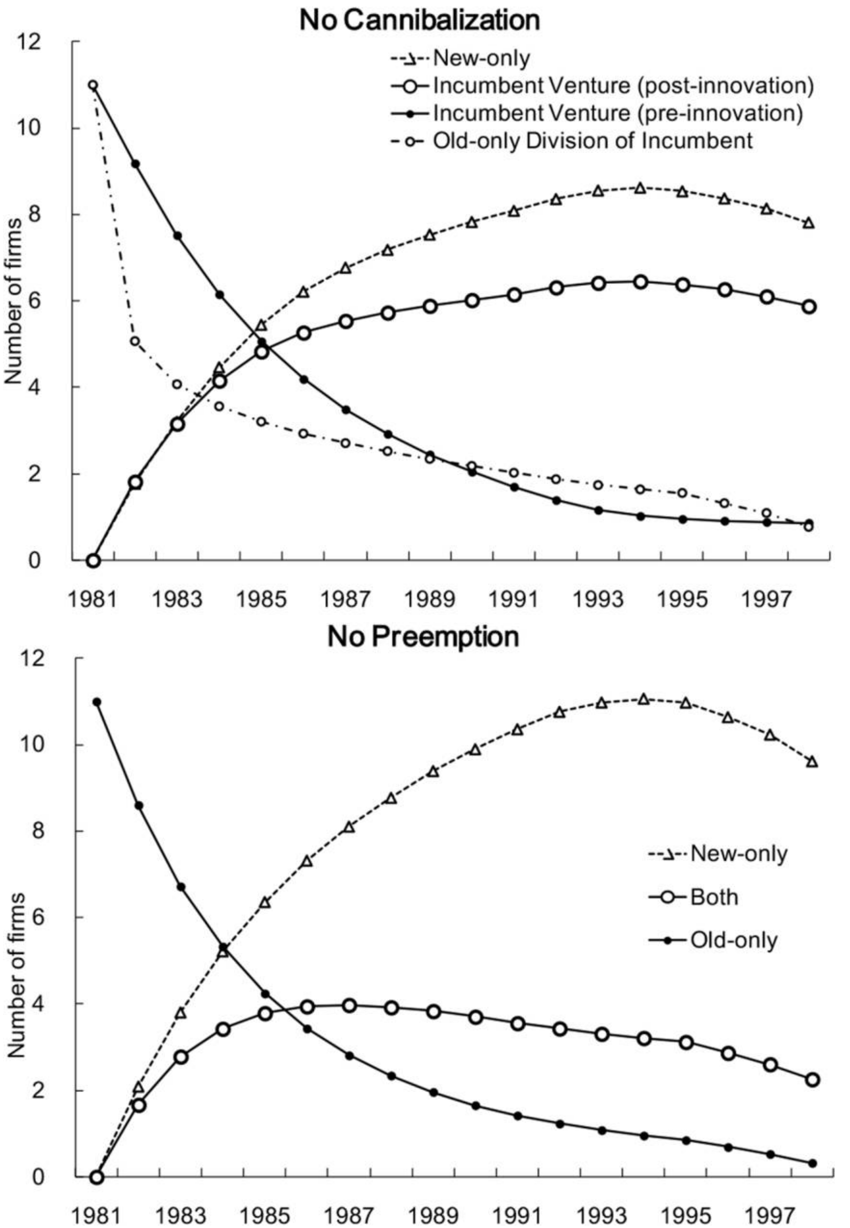

FIG. 6.—Explaining the innovator's dilemma (1). The mean numbers of firms across 10,000 counterfactual simulations.

of their own innovations.[46] By contrast, the no-preemption counterfactual shuts down this channel, thereby reducing the innovation incentives of incumbents.[47]

The simulation results in figure 6 (bottom panel) suggest that without preemptive motives, only 6.02 incumbents would innovate in the first 10 years, compared with 6.30 in the baseline model. The eventual incumbent-entrant innovation gap would widen by 38 percent to 8.91 firms from 6.45 in the baseline model. Thus dynamic strategic interactions appear to be an important driver of industry evolution, even when the market features over a dozen firms.

C. *Heterogeneous Sunk Costs of Innovation*

Third, I quantify the impact of sunk costs by simulating industry dynamics under counterfactual values of κ^{inc}. Although the observed delay of incumbents' innovation relative to entrants would naturally lead us to suspect that incumbents might have faced higher costs of innovation than entrants (i.e., $\kappa^{\text{inc}} > \kappa^{\text{ent}}$), their estimates in the previous section suggest otherwise (i.e., $\kappa^{\text{inc}} < \kappa^{\text{ent}}$). That is, incumbents innovate more slowly than entrants despite their relative advantage in the efficiency of innovation. Therefore, I investigate how much more advantage incumbents would need to innovate as fast as entrants, by solving alternative industry equilibria under lower values of $\kappa^{\text{inc}} < \hat{\kappa}^{\text{inc}} = 1.24$ while holding κ^{ent} fixed at $\hat{\kappa}^{\text{ent}} = 2.25$.

Figure 7 (top panel) evaluates the incumbent-entrant innovation gap at different values of $\kappa^{\text{inc}} \in [0, 1.2]$ and shows that incumbents need practically zero sunk cost to match the pace of entrants' innovations. Figure 7 (bottom panel) visually confirms this finding by showing the evolution of market structure when $\kappa^{\text{inc}} = 0$.

Thus we may summarize the findings of this section as follows: despite strong preemptive motives and innovation cost advantages over entrants, incumbents are reluctant innovators because of cannibalization. For

[46] For example, the estimated model suggests a potential entrant at the end of 1985 who expects that $s_{1986} = (0, 0, 0, 4)$ would enter/innovate 50.5 percent of the time, but this probability drops to 44.5 percent, 41.2 percent, 38.1 percent, and 36.0 percent when N_{1986}^{both} increases to 1, 2, 3, and 4, respectively.

[47] Note that the counterfactual potential entrants follow the baseline strategy, but with a modified state, and are therefore not best-responding to incumbents' innovations. This "optimization error" is by design, because we cannot shut down preemptive motives from incumbents' decision making as long as entrants best-respond to incumbents' strategy. Alternatively, Chicu (2012) used an open-loop equilibrium, which may well shut down "reactive" motives by assuming commitment, in the sense that players do not act on all payoff-relevant states. But that does not mean that incumbents lose incentives to innovate earlier than entrants (i.e., preemptive motives), because it is still a Nash equilibrium (i.e., mutual best responses).

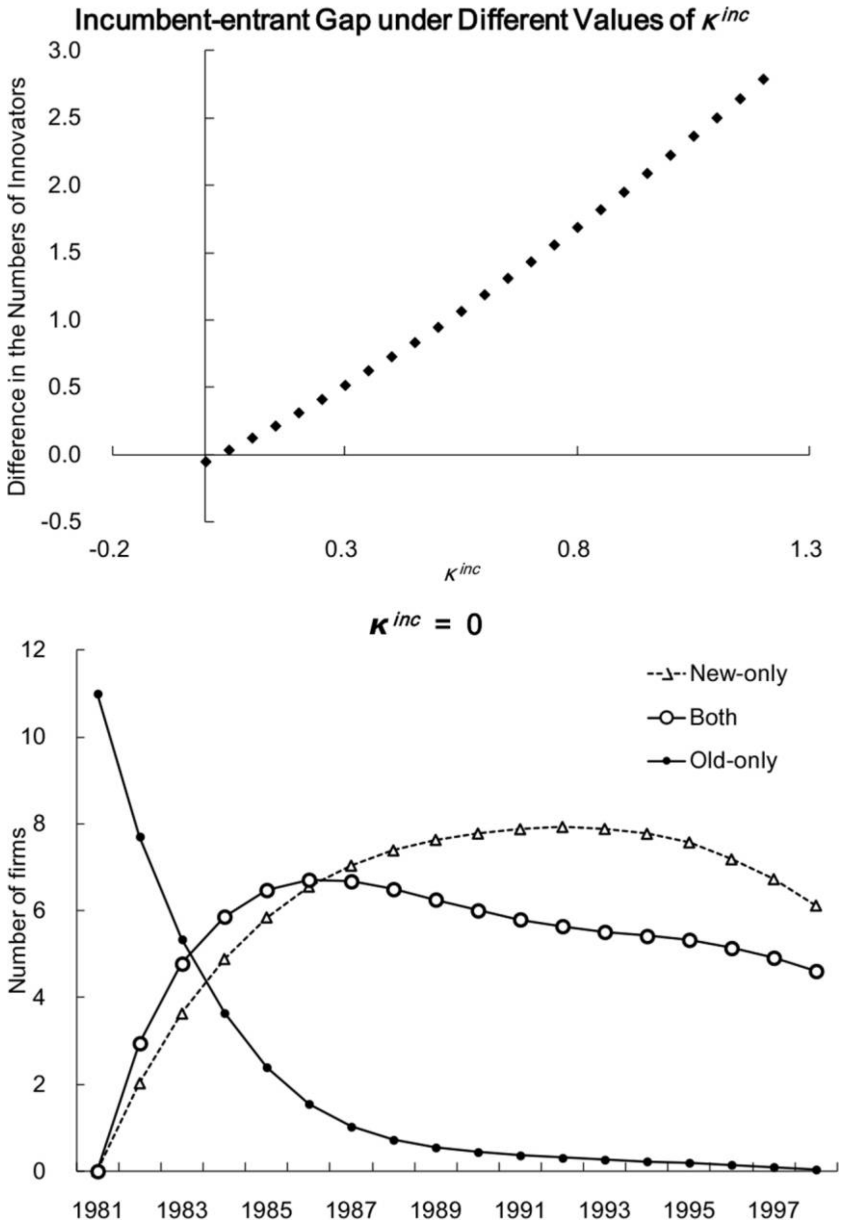

FIG. 7.—Explaining the innovator's dilemma (2). The innovation gap in the top panel refers to the difference in the numbers of innovators among incumbents and entrants during the first 10 years, and a positive gap indicates entrants' lead over incumbents (see eq. [12]). The bottom panel shows the mean numbers of firms across 10,000 counterfactual simulations.

further discussions on another fundamental asymmetry between incumbents and entrants (i.e., the option value of waiting and the difference between internal and external financing), see online appendix A.4.2, in which all of the three theoretical forces are shut down at the same time.

VII. Alternative Hypotheses

This section considers two alternative hypotheses to explain the observed patterns of industry dynamics. One is Christensen's (1997) explanations of the innovator's dilemma, most of which implicitly rely on some notion of bounded rationality and persistent heterogeneity across firms. The other is a version of Jovanovic's (1982) learning and selection model, important components of which are asymmetric information as well as persistent heterogeneity. Data limitations prevent my empirical model from explicitly incorporating these factors, but my estimates will reflect some of these factors if they existed in the true data-generating process. These alternative hypotheses are not necessarily inconsistent with my findings, and we can discuss their implications within the framework of this paper to some extent.

A. Christensen's (1997) Cognitive and Organizational Biases

Christensen's original work employed a qualitative/descriptive approach and relied primarily on interviews. His work does not contain a formal model, and he entertained various hypotheses including (1) incumbents' myopia, (2) incumbents' cognitive biases against the new products, (3) the internal resource struggle between the existing and the new business teams, and (4) the influence of the existing customers. The first hypothesis suggests that the managers of incumbent firms emphasized static profit maximization with respect to the existing products rather than forward-looking investments in the new products. The second hypothesis suggests that the senior management of incumbent firms dismissed the new products as irrelevant. The third hypothesis suggests that the old-product group won most budgets and engineers at the expense of the new-product group. The fourth hypothesis suggests that the buyers of the old products contributed disproportionately to the incumbents' current profits, and hence the latter's top management prioritized the existing business. These stories appear to be realistic descriptions of what could go wrong within established firms with respect to the introduction of new technologies.

Heterogeneous beliefs across managers and the internal organization of the firm are outside the scope of my model. Nevertheless, my estimates of the sunk costs of innovation, κ^{inc} and κ^{ent}, will reflect any systematic heterogeneity (with respect to the innovation decisions) between incumbents and entrants in the data. That is, although I do not explicitly model psychological or organizational biases, any systematic difference between incumbents and entrants will manifest itself in the "residuals" of the model, because my empirical analysis employs the principle of revealed preference instead of stated preference (e.g., interviews). The estimates (i.e., $\kappa^{\text{inc}} < \kappa^{\text{ent}}$) suggest that incumbents were reluctant innovators despite net innovation cost advantages over entrants. Thus incumbents seem to have enjoyed sufficiently large (gross) advantages to offset any negative influence of cognitive or organizational biases. See Section V.C for further discussions on the interpretation of κ^{inc} and κ^{ent}.

B. *Jovanovic's (1982) Learning and Selection*

One of the classic models of industry dynamics is Jovanovic's (1982) learning and selection model, in which atomistic (or monopolistically competitive) firms are born with persistent heterogeneity in productivity. He assumes that the firm does not know its true productivity level and has to passively learn about it from the history of its own product market performance (i.e., the sequence of realizations of noisy signals about its innate productivity) by staying in the market. After some length of time, less productive firms will optimally choose to exit the industry, and hence learning and self-selection become the driving force of firm dynamics.

Unlike Hopenhayn's (1992) or Ericson and Pakes's (1995) model, Jovanovic's model does not feature active investment (e.g., innovation) by incumbents, but we may consider a version of this model to theorize about the difference between incumbents and entrants. More specifically, if we are willing to abstract from incumbents' innovation and strategic interactions, we may interpret incumbents and entrants in the current empirical context as "older" and "younger" firms in Jovanovic's model, respectively. On the basis of this interpretation, the model offers the following predictions: (1) incumbents are more productive than entrants, (2) incumbents exit less often than entrants, and (3) incumbents' product market performance and survival do not depend on the emergence of the new technology or the entry of new rivals.

Let us consider the applicability of these predictions to my data. First, incumbents in the data tend to be larger than entrants in a given year, because the demand for 5.25-inch HDDs constituted the lion's share of the entire HDD sales during the first half of the sample period, and innovating incumbents sold both 5.25- and 3.5-inch HDDs during the second half. If we assume that the firm size is an increasing function of productivity, the model's first prediction fits this data pattern.

Second, evidence is mixed regarding exit rates. The same number of incumbent innovators and actual entrants existed in 1985, and we may compare their survival rates 10 years later (i.e., after the shakeout). Of the seven incumbent innovators in 1985, four were still active in 1995. Likewise, three of the seven actual entrants in 1985 survived until 1995. Thus more incumbents of the 1985 cohort survived than entrants of the same cohort, but the difference does not seem definitive.[48]

Third, the market for 5.25-inch HDDs all but disappeared by the end of the sample period, and no major firm survived as a producer of 5.25-inch HDDs alone. Moreover, of the 11 major incumbents in the data, seven exited during the sample period, and their exit timing is concentrated between 1989 and 1993, which coincides with the shift in demand from 5.25- to 3.5-inch HDDs and lags the massive entry of new rivals (i.e., most entry happened before 1988). HDDs are standardized goods that are traded globally, with concentrated market structure and limited room for brand- or location-based differentiation. Thus the nonstrategic model's prediction appears less useful with regard to competitive interactions in this industry.

Now let us consider how my model incorporates these features. First, incumbents in my model share the same productivity level as entrants (in terms of the marginal cost of producing 3.5-inch HDDs) but are larger than entrants because 5.25-inch HDDs were important in the early years, and innovating incumbents sold both 5.25- and 3.5-inch HDDs in the later years. Second, incumbents in my model exit slightly less often than entrants, because $\pi_t^{\text{both}} = \pi_t^{\text{old}} + \pi_t^{\text{new}} \geq \pi_t^{\text{new}}$ for all t (see fig. 4), which implies $V_t^{\text{both}} \geq V_t^{\text{new}}$, and hence $pr_t^{\text{both}}(d_{it} = \text{exit}) \leq pr_t^{\text{new}}(d_{it} = \text{exit})$ (see Sec. IV.C). Third, in my model, incumbents' profits and survival depend on both their own innovation decisions and entrants' decisions, which fits the data pattern described above and is the main focus of this paper.

In conclusion, Jovanovic's model shares two qualitatively similar predictions with my model, but its nonstrategic nature limits its applicability to this paper's empirical context in which incumbents and entrants produced similar goods and competed with each other in a concentrated market.

In principle, I can conceive an extension of my model to incorporate passive learning and persistent unobserved heterogeneity. In practice, however, its implementation faces two challenges. First, if the firm has to learn about its own productivity, it should also need to learn about its rivals' productivity, the strategic interactions of which would materially

[48] Seven incumbent innovators and 10 actual entrants existed in 1988. For this cohort, the 10-year survival rates are 57 percent and 30 percent, respectively, which is an outcome that is more consistent with the selection story. Thus the analysis of survivorship is sensitive to the choice of base year.

complicate the model. Second, the identification of a model with informational stories typically requires more detailed data than what is currently available. The key parameters would include the precision of signals, the learning process, and the firm-year-specific productivity, all of which are difficult to identify convincingly with my data. For these reasons, learning and persistent heterogeneity are outside the scope of this paper and represent opportunities for future research.

VIII. Policy Experiments

This section evaluates welfare performance of public policies concerning innovation and competition by conducting counterfactual simulations. Specifically, I experiment with two policies: (1) a broad patent on new HDDs and (2) license fees. Online appendix A.4.3 reports an additional experiment in which I simulate the impact of an international trade dispute over intellectual property rights.

A. Broad Patent

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 the industry (see table 6). After years of lawsuits between Rodime and its rivals, the US Court of Appeals for the Federal Circuit (CAFC), a centralized appellate court for patent cases established in 1982, 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.

Although Rodime's claims were considered outrageous in the industry at the time and ended in a legal gray 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 preannounced broad patent regime and the other to study the impact of ex post "surprise" court rulings.

In the first, ex ante counterfactual, only the first innovators can legally produce and sell new HDDs. The patent authority announces and precommits to this legal arrangement before 1981, when the dynamic oli-

TABLE 6
Brief History of Rodime's 3.5-Inch HDD Patent Affair

Year	Events
1980	Rodime became independent from Burroughs's 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 US 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 it 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 US 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*.

* Key events that motivate my counterfactual simulations. See Igami and Subrahmanyam (2015) for a study of statistical relationships between patents and innovations in the HDD industry.

gopoly game begins.[49] Figure 8 (top panel) shows the industry evolution under the preannounced broad patent regime. On average, 6.7 of the 11 incumbents innovate in 1981 and no new entry occurs, which leads to a more concentrated market structure than in the baseline case;[50] so consumer

[49] Operationally, the game now contains an additional state variable that indicates whether some firms have already innovated, in which case the remaining pre-innovation incumbents can no longer choose to innovate and potential entrants cannot enter. All of these rules are common knowledge from 1981.

[50] The incumbents' incentive to delay innovation continues to exist in this simulation. However, the new legal setting makes preemptive motives extremely strong, completely dominating the cannibalization concern. My model and data do not distinguish between invention, patenting, and commercialization, but I found that a monopolist (without any active or potential rival) would choose to innovate in 1986 in one of my unreported simulation results. Thus we may deduce that the first innovators/patent holders in this counterfactual would optimally delay the actual sale of new HDDs until 1986, 5 years after filing for patent protection, if they could perfectly collude.

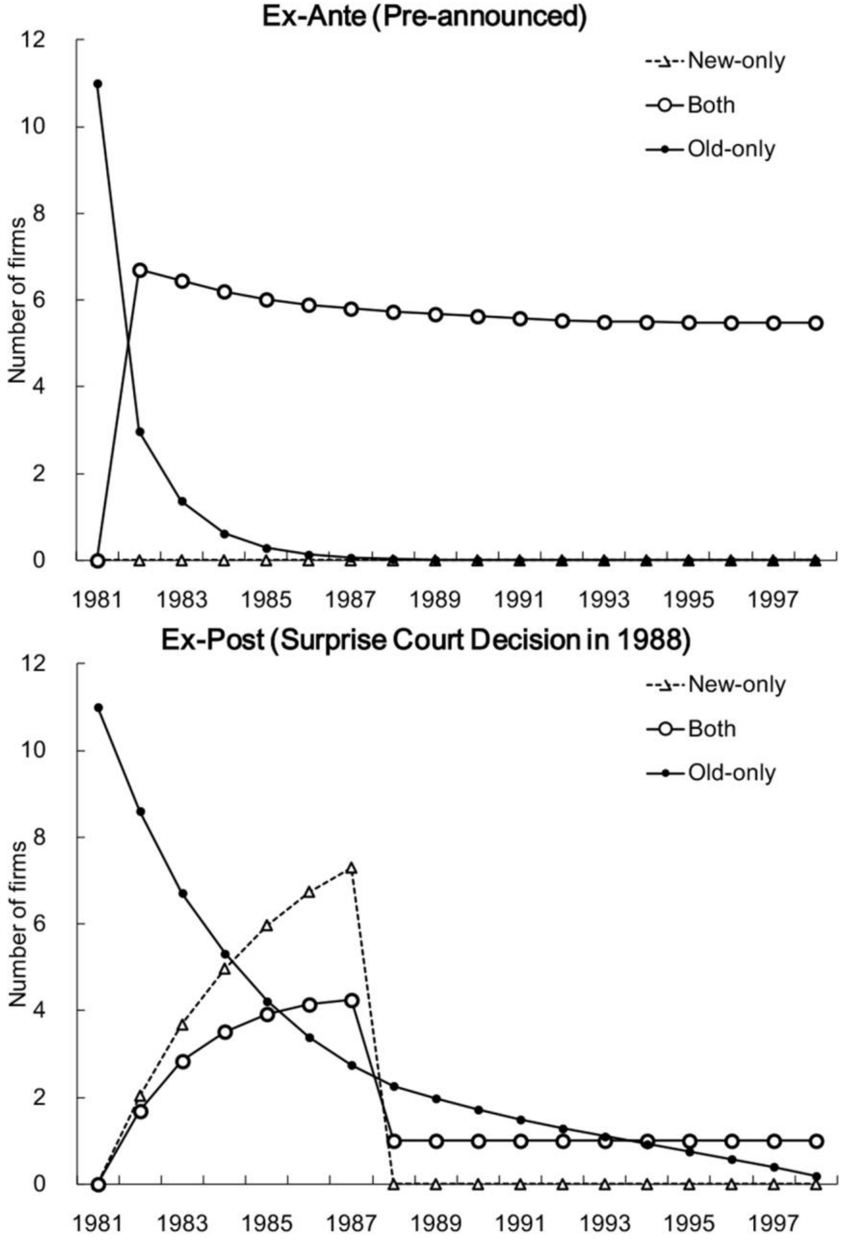

FIG. 8.—Policy experiment 1 (broad patent): the mean numbers of firms across 10,000 counterfactual simulations.

TABLE 7
Welfare Analysis of Innovation/Competition Policies

	Consumer Surplus (1)	Producer Surplus (2)	Fixed Cost of Operation (3)	Sunk Cost of Innovation (4)	Social Welfare (1 + 2 + 3 + 4)	Change from Baseline
Baseline	52.7	4.0	−9.7	−25.7	21.3	. . .
Broad patent:						
Ex ante	42.1	6.9	−5.7	−8.7	34.6	+62.5%
Ex post	28.3	4.9	−8.0	−23.2	2.0	−90.8%
License fees:						
25%	51.2	4.8	−9.6	−25.0	21.5	+.7%
50%	50.5	5.2	−9.6	−24.1	22.0	+3.3%
75%	47.2	6.5	−9.4	−24.1	20.1	−5.8%

Note.—Each number is the sum of discounted present values as of 1981.

surplus drops by 20.1 percent to \$42.1 billion and producer surplus increases by 72.5 percent to \$6.9 billion (table 7), highlighting the anticompetitive effect of patent protection. This deadweight loss, however, should be judged against the cost savings thanks to the reduction of excess entry, which is sizable and leads to a net gain of 62.5 percent in social welfare. Because incumbents turned out to be more efficient innovators (i.e., $\kappa^{\text{inc}} < \kappa^{\text{ent}}$), preventing the duplication of efforts by entrants is effective in saving societywide sunk costs. Moreover, the innovation timing is front-loaded to 1981, which is a favorable development from the viewpoint of pro-innovation policy. Unfortunately, such a well-defined patent regime is probably unrealistic in the computer-related industries, in which technologies are complex and both the authorities and firms suffer from the lack of predictability; but this counterfactual experiment would suggest ample room for social gains if such policies were feasible.

By contrast, the second simulation focuses on a more realistic, ex post enforcement of patents. In this counterfactual scenario, Rodime's rivals ignore the company's patent claims until 1988, when the CAFC announces its surprise ruling to honor Rodime's patent infringement claims, paving the way for a legal monopoly of the 3.5-inch technology. As a result of the preliminary injunction to stop infringing activities, the other producers of new HDDs immediately go out of business.[51] In stark contrast with the previous experiment, social welfare drops by 90.8 percent in this scenario because most firms had already paid the sunk costs and started production of 3.5-inch HDDs by 1988, so that no cost savings exist. Instead, the only major change is that the industry becomes a quasi monopoly and consumers thus suffer. These developments represent thought experiments of an extreme case, of course, and patent lawsuits in reality usually end with

[51] Operationally, the industry evolves just as in the baseline model until 1988, because firms do not foresee the CAFC's surprise ruling. Then, from 1988, market structure permanently changes to $N^{\text{both}} = 1$ and $N^{\text{new}} = N^{\text{pe}} = 0$, and the remaining old-only firms may no longer innovate.

settlements and licensing agreements (see below). Nevertheless, I believe that this counterfactual contains a grain of truth in that (1) patent disputes arise well after innovation costs are incurred, so that no cost savings exist; and (2) firms are willing to settle and pay license fees to "patent trolls" even when the latter's infringement claims seem bogus, precisely because managers wish to eliminate the possibility of destructive rare events such as the preliminary injunction to stop operations. Thus I interpret this result as a cautionary note indicating an upper bound of how poorly a dysfunctional patent policy could perform.

B. License Fees

The broad-patent counterfactuals assume that the 3.5-inch HDD patents prevent all rival firms from using the new technology, which is an extreme form of enforcing intellectual property rights. In practice, patents usually lead to license fees. The licensor may earn license fees from its rivals, but the size of the reward would not be as large as the hypothetical monopoly profit (simulated in Sec. VII.A). Likewise, its rivals may become licensees and transfer part of their profits to the patent owner, which would reduce their profitability and enterprise values, but they would not be forced to exit by the preliminary injunction. Thus simulating a broad patent regime with license fees will help us establish the likely impact of policy interventions in a more subtle manner.

For these purposes, I conduct ex post patent counterfactuals with three different rates of license fees (25 percent, 50 percent, and 75 percent). Table 7 (bottom rows) shows that the welfare consequences become more nuanced, with net changes in social welfare at +0.7 percent, +3.3 percent, and −5.8 percent, respectively. The underlying economic mechanism is as follows. Although the change in market structure is not as drastic as in the legal monopoly case of Section VII.A, license fees also discourage innovation/entry (and encourage exit) after 1988, leading to more concentrated market structures than in the baseline model (see fig. 9). As a result of reduced competition, consumer surplus decreases between 2.8 percent and 10.5 percent whereas producer surplus increases between 21.9 percent and 63.5 percent (most of which accrues to Rodime as the sole licensor). At the same time, the reduced number of firms after 1988 also implies slightly lower fixed costs and sunk costs at the aggregate level, partially offsetting the deadweight losses discussed above. This trade-off between deadweight loss and excess entry is so subtle that the net welfare impact of license fees could be either positive or negative, depending on the specific rate of fees.[52]

[52] Another important public policy parameter is the timing of the hypothetical surprise ruling by the CAFC, which is held fixed in 1988 across my simulations. In principle, the later the timing of intervention, the worse its welfare impact becomes, because the social cost savings from reduced entry/innovation will be less.

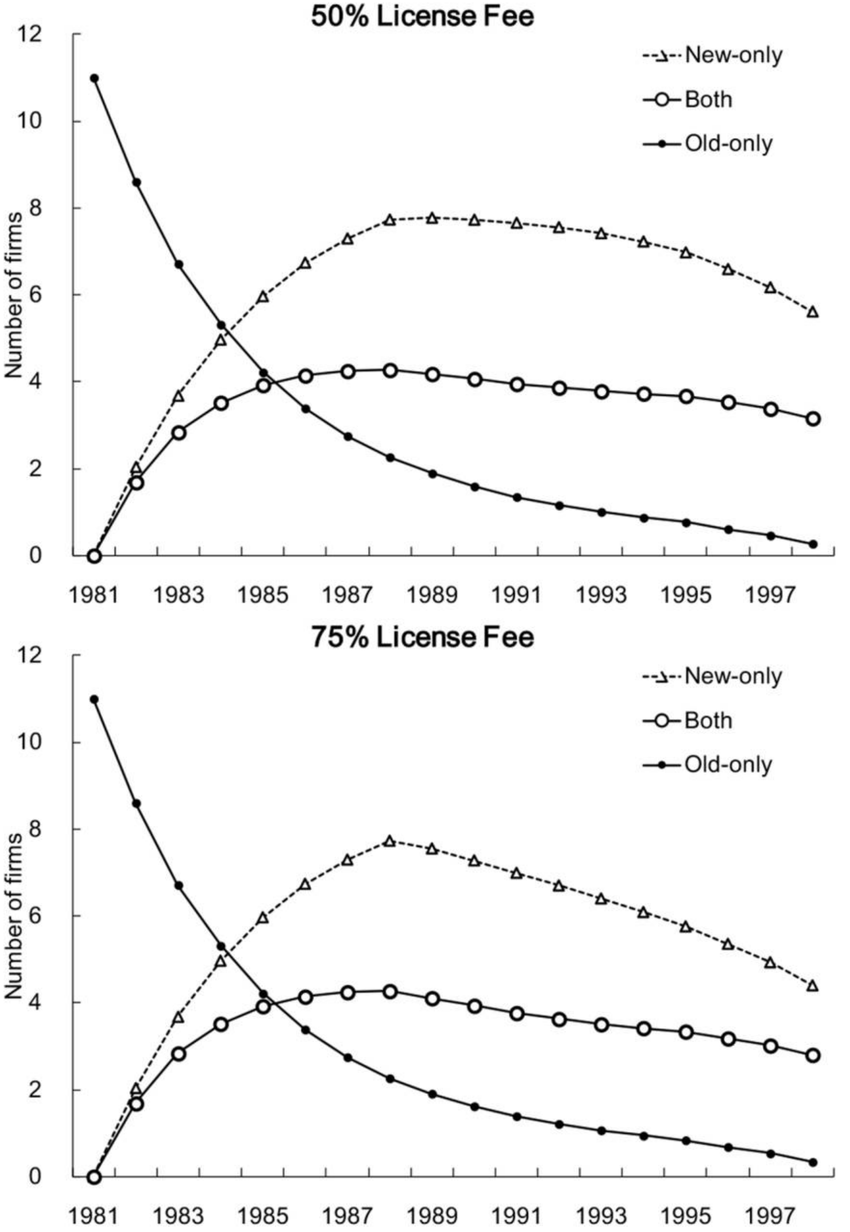

FIG. 9.—Policy experiment 2 (license fees): the mean numbers of firms across 10,000 counterfactual simulations.

IX. Conclusion

I analyzed the strategic industry dynamics of creative destruction in this paper. Although commonly understood as a turnover of technologies alone, an important implication of creative destruction hinges on the simultaneous turnovers of technologies and firms, because Schumpeter (1942) proposed this concept as an answer to his own historical question about market power.[53] The results from my empirical analysis of the HDD industry suggest that despite strong preemptive motives and a higher efficiency of innovation than entrants, cannibalization makes incumbents reluctant to innovate, which is a product innovation analogue of Arrow's (1962) replacement effect. In other words, a systematic reason exists for incumbents to delay innovations, even if they are rational and do not suffer from informational or organizational disadvantages. These results are specific to the HDD market, but the economic incentives studied here are general and could be expected to operate in other innovative industries, many of which are oligopolistic.[54]

Counterfactual experiments suggest substantial welfare gains from a well-defined patent system if it worked properly as an ex ante incentive scheme. However, such a policy is probably infeasible in a complex technology space, such as computer-related technologies. A more realistic ex post granting of monopoly rights exhibits disappointing welfare performance, and so does the license fee–based implementation of such a patent regime. The failure of these typical "pro-innovation" government interventions might appear to be negative findings, but the flip side of the results is that the HDD industry performed quite well in the actual course of history. And this finding resonates with Schumpeter's conjecture that sufficient competition and innovation exist in the long-run process of creative destruction.

References

Acemoglu, Daron, and Dan Cao. 2010. "Innovation by Entrants and Incumbents." Manuscript, Massachusetts Inst. Tech.

Acs, Zoltan J., and David B. Audretsch. 1988. "Innovation in Large and Small Firms: An Empirical Analysis." *A.E.R.* 78 (4): 678–90.

Aguirregabiria, Victor, and Chun-Yu Ho. 2012. "A Dynamic Oligopoly Game of the US Airline Industry: Estimation and Policy Experiments." *J. Econometrics* 168 (1): 156–73.

[53] Schumpeter (1942) asked why the aggregate output could grow and the standard of life improve so persistently between the 1890s and the 1940s, which in his view represented "the period of relatively unfettered 'big business'" and would usually imply reduced production and higher prices in a static model. His answer was that the competition from the new commodity and the new technology would enforce more competitive behavior in the long run.

[54] That innovative industries tend to be global oligopolies is not a coincidence. See Sutton (1998) for detailed theoretical explanations.

Aguirregabiria, Victor, and Pedro Mira. 2007. "Sequential Estimation of Dynamic Discrete Games." *Econometrica* 75 (1): 1–53.

Aguirregabiria, Victor, and Junichi Suzuki. 2014. "Identification and Counterfactuals in Dynamic Models of Market Entry and Exit." *Quantitative Marketing and Econ.* 12 (3): 267–304.

Arrow, Kenneth J. 1962. "Economic Welfare and the Allocation of Resources to Invention." In *The Rate and Direction of Economic Activity*, edited by R. R. Nelson. Princeton, NJ: Princeton Univ. Press.

———. 1974. *The Limits of Organization.* New York: Norton.

Aw, Bee Yan, Mark J. Roberts, and Daniel Yi Xu. 2009. "R&D Investment, Exporting, and Productivity Dynamics." Working Paper no. 14670, NBER, Cambridge, MA.

Bajari, Patrick, C. Lanier Benkard, and Jonathan Levin. 2007. "Estimating Dynamic Models of Imperfect Competition." *Econometrica* 75 (5): 1331–70.

Benkard, C. Lanier. 2004. "A Dynamic Analysis of the Market for Wide-Bodied Commercial Aircraft." *Rev. Econ. Studies* 71:581–611.

Berry, Steven T. 1994. "Estimating Discrete-Choice Models of Product Differentiation." *RAND J. Econ.* 25 (2): 242–62.

Berry, Steven T., and Philip A. Haile. 2009. "Nonparametric Identification of Multinomial Choice Demand Models with Heterogeneous Consumers." Working Paper no. 15276, NBER, Cambridge, MA.

Berry, Steven T., James Levinsohn, and Ariel Pakes. 1995. "Automobile Prices in Market Equilibrium." *Econometrica* 63:841–90.

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

———. 2003. "Pro-Innovation Competition Policy: Microsoft and Beyond." Discussion paper, Competition Policy Res. Center, Fair Trade Comm. Japan, Tokyo.

Chesbrough, Henry. 1999. "Arrested Development: The Experience of European Hard Disk Drive Firms in Comparison to U.S. and Japanese Firms." *J. Evolutionary Econ.* 9:287–330.

Chicu, Mark. 2012. "Dynamic Investment and Deterrence in the U.S. Cement Industry." Manuscript, Northwestern Univ.

Christensen, Clayton M. 1993. "The Rigid Disk Drive Industry: A History of Commercial and Technological Turbulence." *Bus. Hist. Rev.* 67:531–88.

———. 1997. *The Innovator's Dilemma.* New York: HarperBusiness.

Cohen, Wesley M. 2010. "Fifty Years of Empirical Studies of Innovative Activity and Performance." In *Handbook of the Economics of Innovation*, vol. 1, edited by Bronwyn H. Hall and Nathan Rosenberg. Amsterdam: Elsevier.

Collard-Wexler, Allan. 2013. "Demand Fluctuations in the Ready-Mix Concrete Industry." *Econometrica* 81 (3): 1003–37.

Dixit, Avinash K., and Robert S. Pindyck. 1994. *Investment under Uncertainty.* Princeton, NJ: Princeton Univ. Press.

Egesdal, Michael, Zhenyu Lai, and Che-Lin Su. 2014. "Estimating Dynamic Discrete-Choice Games of Incomplete Information." Manuscript, Booth School Bus., Univ. Chicago.

Ehrnberg, Ellinor, and Niklas Sjöberg. 1995. "Technological Discontinuities, Competition and Firm Performance." *Technology Analysis and Strategic Management* 7 (1): 93–108.

Ericson, Richard, and Ariel Pakes. 1995. "Markov-Perfect Industry Dynamics: A Framework for Empirical Work." *Rev. Econ. Studies* 62 (1): 53–82.

Franco, April Mitchell, and Darren Filson. 2006. "Spin-Outs: Knowledge Diffusion through Employee Mobility." *RAND J. Econ.* 37 (4): 841–60.

Fudenberg, Drew, and Jean Tirole. 1986. *Dynamic Models of Oligopoly*. London: Routledge.

Futures Group. 1984. *Characterization of Innovations Introduced on the U.S. Market in 1982*. Nat. Tech. Information Service Document PB-84-212067. Washington, DC: US Small Bus. Admin., Office of Advocacy.

Gellman Research Associates. 1976. *Indicators of International Trends in Technological Innovation*. Final Report, National Science Foundation, Nat. Tech. Information Service Document PB-263-738. Jenkintown, PA: Gellman Res. Assoc.

———. 1982. *The Relationship between Industrial Concentration, Firm Size and Technological Innovation*. Final Report, US Small Bus. Admin. Glastonbury, CT: Futures Group.

Gilbert, Richard. 2006. "Looking for Mr. Schumpeter: Where Are We in the Competition-Innovation Debate?" In *Innovation Policy and the Economy*, vol. 6, edited by Adam B. Jaffe, Josh Lerner, and Scott Stern. Cambridge, MA: MIT Press.

Gilbert, Richard, and David Newbery. 1982. "Preemptive Patenting and the Persistence of Monopoly." *A.E.R.* 72 (2): 514–26.

Goettler, Ronald, and Brett Gordon. 2011. "Does AMD Spur Intel to Innovate More?" *J.P.E.* 119 (6): 1141–1200.

Gort, Michael, and Steven Klepper. 1982. "Time Paths in the Diffusion of Product Innovations." *Econ. J.* 92:630–53.

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

Hall, Bronwyn. 2004. "Innovation and Diffusion." In *Oxford Handbook of Innovation*, edited by Jan Fagerberg, David Mowery, and Richard Nelson. Oxford: Oxford Univ. Press.

Hashmi, Aamir Rafique, and Johannes Van Biesebroeck. 2016. "The Relationship between Market Structure and Innovation in Industry Equilibrium: A Case Study of the Global Automobile Industry." *Rev. Econ. and Statis.* 98 (1): 192–208.

Hausman, Jerry. 1996. "Valuation of New Goods under Perfect and Imperfect Competition." In *The Economics of New Goods*, edited by T. Bresnahan and R. Gordon. Studies in Income and Wealth, vol. 58. Chicago: Univ. Chicago Press (for NBER).

Henderson, Rebecca M. 1993. "Underinvestment and Incompetence as Responses to Radical Innovation: Evidence from the Photolithographic Alignment Industry." *RAND J. Econ.* 24 (2): 248–70.

Henderson, Rebecca M., and Kim B. Clark. 1990. "Architectural Innovation: The Reconfiguration of Existing Product Technologies and the Failure of Established Firms." *Administrative Sci. Q.* 35:9–30.

Hopenhayn, Hugo A. 1992. "Exit, Selection, and the Value of Firms." *J. Econ. Dynamics and Control* 16 (3–4): 621–53.

Igami, Mitsuru. Forthcoming. "Industry Dynamics of Offshoring: The Case of Hard Disk Drives." *American Econ. J.: Microeconomics*.

Igami, Mitsuru, and Jai Subrahmanyam. 2015. "Patent Statistics as an Innovation Indicator? Evidence from the Hard Disk Drive Industry." Manuscript, Yale Univ.

Igami, Mitsuru, and Kosuke Uetake. 2015. "Mergers, Innovation, and Entry-Exit Dynamics: The Consolidation of the Hard Disk Drive Industry, 1996–2015." Manuscript, Yale Univ.

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 Univ. Press.

Jovanovic, Boyan. 1982. "Selection and the Evolution of Industry." *Econometrica* 50 (3): 649–70.

Kim, Myongjin. 2013. "Strategic Responses to Used Goods Markets: Airbus and Boeing." Manuscript, Univ. Oklahoma.

King, Andrew A., and Christopher L. Tucci. 2002. "Incumbent Entry into New Market Niches: The Role of Experience and Managerial Choice in the Creation of Dynamic Capabilities." *Management Sci.* 48:171–86.

Klepper, Steven. 1996. "Entry, Exit, Growth, and Innovation over the Product Life Cycle." *A.E.R.* 86 (3): 562–83.

———. 2002. "Firm Survival and the Evolution of Oligopoly." *RAND J. Econ.* 33 (1): 37–61.

Klepper, Steven, and Elizabeth Graddy. 1990. "The Evolution of New Industries and the Determinants of Market Structure." *RAND J. Econ.* 21 (1): 27–44.

Klepper, Steven, and Kenneth L. Simons. 2000. "The Making of an Oligopoly: Firm Survival and Technological Change in the Evolution of the U.S. Tire Industry." *J.P.E.* 108 (4): 728–60.

Klette, Tor Jakob, and Samuel Kortum. 2004. "Innovating Firms and Aggregate Innovation." *J.P.E.* 112 (5): 986–1018.

Kreps, David M., and Jose A. Scheinkman. 1983. "Quantity Precommitment and Bertrand Competition Yield Cournot Outcomes." *Bell J. Econ.* 14 (2): 326–37.

Lentz, Rasmus, and Dale T. Mortensen. 2008. "An Empirical Model of Growth through Product Innovation." *Econometrica* 76 (6): 1317–73.

Lerner, Josh. 1997. "An Empirical Exploration of a Technology Race." *RAND J. Econ.* 28 (2): 228–47.

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 Univ. Press.

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

Pakes, Ariel, Michael Ostrovsky, and Steven Berry. 2007. "Simple Estimators for the Parameters of Discrete Dynamic Games (with Entry/Exit Examples)." *RAND J. Econ.* 38 (2): 373–99.

Pavit, K., M. Robson, and J. Townsend. 1987. "The Size Distribution of Innovating Firms in the UK: 1945–1983." *J. Indus. Econ.* 35 (3): 297–316.

Pesendorfer, Martin, and Philipp Schmidt-Dengler. 2008. "Asymptotic Least Squares Estimators for Dynamic Games." *Rev. Econ. Studies* 75 (3): 901–28.

Rust, John. 1987. "Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher." *Econometrica* 55 (5): 999–1033.

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

Sah, R. J. and J. E. Stiglitz. 1986. "The Architecture of Economic Systems: Hierarchies and Polyarchies." *A.E.R.* 76:716–27.

Scherer, F. M. 1965. "Firm Size, Market Structure, Opportunity, and the Output of Patented Inventions." *A.E.R.* 55 (5): 1097–1125.

Scherer, F. M., and D. Ross. 1990. *Industrial Market Structure and Economic Performance*. Boston: Houghton Mifflin.

Schmidt-Dengler, Philipp. 2006. "The Timing of New Technology Adoption: The Case of MRI." Manuscript, London School Econ.

Schumpeter, Joseph A. 1934. *The Theory of Economic Development*. Cambridge, MA: Harvard Univ. Press.

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

Seim, Katja. 2006. "An Empirical Model of Firm Entry with Endogenous Product-Type Choices." *RAND J. Econ.* 37 (3): 619–40.

Stoneman, Paul, and Giuliana Battisti. 2010. "The Diffusion of New Technology." In *Handbook of the Economics of Innovation*, vol. 1, edited by Bronwyn H. Hall and Nathan Rosenberg. Amsterdam: Elsevier.

Su, Che-Lin. 2014. "Estimating Discrete-Choice Games of Incomplete Information: Simple Static Examples." *Quantitative Marketing and Econ.* 12 (2): 167–207.

Sutton, John. 1991. *Sunk Costs and Market Structure: Price Competition, Advertising, and the Evolution of Concentration.* Cambridge, MA: MIT Press.

———. 1998. *Technology and Market Structure: Theory and History.* Cambridge, MA: MIT Press.

———. 2013. *Competing in Capabilities: The Globalization Process.* Oxford: Oxford Univ. Press.

Sweeting, Andrew. 2013. "Dynamic Product Repositioning in Differentiated Product Markets: The Effect of Fees for Musical Performance Rights on the Commercial Radio Industry." *Econometrica* 81 (5): 1763–1803.

Tripsas, Mary. 1997. "Surviving Radical Technological Change through Dynamic Capability: Evidence from the Typesetter Industry." *Indus. and Corporate Change* 6 (2): 341–77.

Tushman, Michael L., and Philip Anderson. 1986. "Technological Discontinuities and Organizational Environments." *Administrative Sci. Q.* 31 (3): 439–65.

Xu, Daniel Yi. 2008. "A Structural Model of R&D, Firm Heterogeneity, and Industry Evolution." Manuscript, New York Univ.