Mergers, Innovation, and Entry-Exit Dynamics: Consolidation of the Hard Disk Drive Industry, 1996–2015*

Mitsuru Igami† Kosuke Uetake‡

July 8, 2016

Abstract

How far should an industry be allowed to consolidate when competition and innovation are endogenous? We extend Rust's (1987) framework to incorporate a stochastically alternating-move game of dynamic oligopoly, and estimate it using data from the hard disk drive industry, in which a dozen global players consolidated into only three in the last 20 years. We find plateaux in the estimated equilibrium competition-innovation relationships, with systematic heterogeneity across time and productivity. Our counterfactual simulations suggest the optimal policy should stop mergers when five or less firms exist, highlighting a dynamic welfare tradeoff between ex-post pro-competitive effects and ex-ante value-destruction side effects.

Keywords: Antitrust, Competition and innovation, Dynamic oligopoly, Dynamic welfare tradeoff, Entry and exit, Horizontal mergers, Industry consolidation.

JEL classifications: L13, L41, L63, O31.

*For comments, we thank John Rust, Paul Ellickson, April Franco, and Joshua Gans, who discussed earlier versions of the paper, as well as participants at various seminars and conferences. For inside stories and insights, we thank Jeff Burke, Tu Chen, Finis Conner, MyungChan Jeong, Peter Knight, Currie Munce, Reggie Murray, Orie Shelef, Lawrence Wu, as well as Mark Geenen and his team at *TRENDFOCUS*, including John Chen, Don Jeanette, John Kim, Tim Luehmann. Financial support from the Yale Center for Customer Insights is gratefully acknowledged.

†Yale Department of Economics. E-mail: mitsuru.igami@yale.edu.

‡Yale School of Management. Email: kosuke.uetake@yale.edu.

1 Introduction

How far should an industry be allowed to consolidate? This is a foundational question for antitrust policy since its inception in 1890 as a countermeasure to merger waves (c.f., Lamoreaux 1985). Conventional merger analysis takes a proposed merger as given and focuses on its immediate effects on competition, which is expected to decrease after a target firm exits, and efficiency, which might increase if sufficient "synergies" materialize.[1] Such a static analysis would be appropriate if mergers were completely random events in isolation from competition and innovation, and if market structure and firms' productivity evolved exogenously over time. However, Demsetz (1973) cautioned that monopolies are often endogenous outcomes of competition and innovation, and Berry and Pakes (1993) conjectured that such dynamic factors could dominate static factors. Indeed, in 100% of high-tech merger cases, the antitrust authority has tried to assess potential impacts on innovation but found little guidance in the economics literature.[2] This paper proposes a tractable dynamic oligopoly model in which mergers, innovation, and entry-exit are endogenous, estimates it using data from the process of industry consolidation among the manufacturers of hard disk drives (HDDs) between 1996 and 2015, and quantifies a dynamic welfare tradeoff by simulating hypothetical merger policies.

Mergers in innovative industries represent an opportunity to kill competition and acquire talents, which make them strategic and forward-looking choices of firms.[3] Besides the static tradeoff between market power and efficiency, merger policy needs to consider both ex-post and ex-ante impact. Ex post, a merger reduces the number of competitors and alters their productivity profile, which will change the remaining firms' incentives for subsequent mergers and innovation. Theory predicts mergers are strategic complements (e.g., Qiu and Zhou 2007), hence a given merger increases the likelihood of subsequent mergers. Its impact on subsequent innovation is more complicated because the competition-innovation relationship crucially hinges on the parameters of demand, supply, and investment (e.g., Sutton 1998). These are the ex-post effects of mergers. These changes in competition and innovation will have ex-ante impacts as well, because a tougher antitrust regime will lower firms' expected profits and option values of staying in the market, which will in turn reduce their ex-ante

[1] See Williamson (1968), Werden and Froeb (1994), and Nevo (2000), for example.

[2] See survey by Gilbert and Greene (2015).

[3] According to Reggie Murray, the founder of Ministor, "Most mergers were to kill competitors, because it's cheaper to buy them than to compete with them. Maxtor's Mike Kennan said, 'We'd rather buy them than have them take us out,' referring to Maxtor's acquisition of Quantum in 2001." (January 22, 2015, in Sunnyvale, CA). See Appendix A for a full list of interviews with industry veterans.

incentives to invest in productivity, survival, and market entry. Thus merger policy faces a tradeoff between the ex-post pro-competitive effects and the ex-ante value-destruction side effects. Their exact balance depends on the parameters of demand, cost, and investment functions, hence the quest for optimal merger policy is a theoretical as well as empirical endeavor.

Three challenges haunt the empirical analysis of merger dynamics in the high-tech context. First, mergers in a concentrated industry are rare events by definition, and the nature of the subject precludes the use of experimental methods, hence a model has to complement sparse data. Second, an innovative industry operates in a nonstationary environment and tends to feature globally concentrated market structure,[4] which creates a methodological problem for the application of two-step estimation approaches, because (at most) only one data point exists in each state of the world, which is too few for nonparametric estimation of conditional choice probabilities (CCPs). Third, workhorse models of dynamic oligopoly games such as Ericson and Pakes (1995) entail multiple equilibria, which preclude the application of full-solution estimation methods such as Rust (1987) because parameter estimates will be inconsistent when a single vector of parameter values predicts multiple strategies and outcomes. We solve these problems by developing a tractable model with unique equilibrium, incorporating the nonstationary environment of the HDD industry, and extending Rust's framework to a dynamic game with stochastically alternating moves.

The paper is organized as follows. In Section 2, we introduce a simple model of a dynamic oligopoly with endogenous mergers, innovation, and entry-exit. We depart from the simultaneous-move tradition of the literature and adopt sequential or alternating moves. An unsatisfactory feature of a sequential-move game is that the assumption on the order of moves will generate an artificial early-mover advantage if the order is deterministic (e.g., Gowrisankaran 1995, 1999; Igami 2015, 2016). Instead, we propose a random-mover dynamic game in which the turn to move arrives stochastically. Dynamic games with stochastically alternating moves have been used as a theoretical tool since Baron and Ferejohn (1989) and Okada (1996). Iskhakov, Rust, and Schjerning (2014, 2016) used it to study competition and innovation. We find it useful as an empirical model as well. We combine this random-mover modeling with the HDD market's fundamental feature that the industry is now mature and declining: a finite horizon. With a finite horizon and stochastically alternating moves, we can solve the game for a unique equilibrium by backward induction from the final period, in which profits and values become zero. At most only one firm moves within a period and

[4]Sutton (1998) explains this feature by low transport costs (per value of product) and high sunk costs.

makes a discrete choice between exit, investment in productivity, or merger proposal to one of the rivals. Thus the dynamic game becomes a finite repetition of an effectively single-agent discrete-choice problem. We estimate the sunk costs associated with these discrete alternatives by using Rust's (1987) maximum-likelihood method with the nested fixed-point algorithm.

In Section 3, we describe key features of the HDD industry and the outline of data. This high-tech industry has experienced massive waves of entry, shakeout, and consolidation, providing a suitable context to study the dynamics of mergers and innovation. We explain several product characteristics and institutional background that will inform our subsequent analysis, such as fierce competition among undifferentiated "brands" and an industry-wide technological trend called Kryder's Law (i.e., exogenous technological improvements in areal density).[5] Our dataset consists of three elements. Panel A contains aggregate HDD shipments, HDD price, disk price, and PC shipments, which we will use to estimate demand in Section 4.1. Panel B is firm-level market shares, which we will use to estimate variable costs and period profits in Section 4.2. Panel C records firms' dynamic choices between merger, innovation, and entry-exit, which we will use to estimate sunk costs in Section 4.3.

In Section 4, we take three steps to estimate (i) demand, (ii) variable costs, and (iii) sunk costs, respectively, each of which pairs a model element and a data element as follows. In Section 4.1, we estimate a log-linear demand model from the aggregate sales data in Panel A, treating each gigabytes (GB) as a unit of homogeneous data-storage services. As instruments for prices, we use two cost-shifters: the price of disks (key components of HDDs) and a time trend, both of which reflect Kryder's Law.[6] To control for demand-side dynamics that could arise from the repurchasing cycle of personal computers (PCs), we also include PC shipments as a demand-shifter.

In Section 4.2, we infer the implied marginal cost of each firm in each period from the observed market shares in Panel B, based on the demand estimates in Section 4.1 and a Cournot model (with heterogeneous costs across firms) as a mode of spot-market competition. The firm's first-order condition (FOC) provides a one-to-one mapping from its observed market share to its marginal cost (productivity). Our preferred interpretation of Cournot competition is Kreps and Scheinkman's (1983) model of quantity pre-commitment followed by price competition, given all firms' cost functions (i.e., productivity levels). Effective

[5] Kryder's Law is an engineering regularity that says the recording density (and therefore storage capacity) of HDDs doubles approximately every 12 months, just like Moore's Law, which says the circuit density (and therefore processing speeds) of semiconductor chips doubles every 18 to 24 months

[6] The modeling of Kryder's Law is outside the scope of this paper, and we regard this industry-wide trend as an exogenous technological process that progresses deterministically.

production capacities are highly "perishable" in our high-tech context, because Kryder's Law makes old manufacturing equipment obsolete within a few quarters. Hence our notion of "quantity pre-commitment" is the amount of re-tooling efforts that each firm makes in each quarter, which determines its effective output capacity for that period. Likewise, the real-world counterpart to our notion of cost (productivity) is intangible assets such as the state of tacit knowledge embodied by teams of engineers, rather than durable physical capacities. Our profit-margin estimates strongly correlate with the accounting profit margins in the firms' income statements.

In Section 4.3, we estimate the sunk costs of merger, innovation, and entry, based on the observed choice patterns in Panel C and the benefits of these actions (i.e., streams of period profits) from Section 4.2. Our dynamic discrete-choice model in Section 2 provides a clear mapping from the observed choices and their associated benefits to the implied costs of these choices, which is analogous to the way Cournot FOC mapped output data and demand elasticity into implied costs. For example, if we observe many mergers despite small incremental profits, the model will reconcile these observations by inferring a low cost of merger: revealed preference.[7] Our firm-value estimates match closely with the actual acquisition prices in the historical merger deals.

In Section 4.4, we investigate the equilibrium relationships between innovation, merger, and market structure, based on our estimates of optimal strategies (i.e., conditional choice probabilities, or CCPs, of innovation and merger) from Section 4.3. Three patterns emerge. First, the incentive to innovate increases steeply as the number of firms increases from 1 to 3, reflecting the dynamic preemption motives as in Gilbert and Newbery (1982). This pattern is robust across years and firms of different productivity levels. Second, this competition-innovation relationship becomes heterogeneous and nonmonotonic with more than 3 firms: (i) the innovation rate increases at a decreasing rate monotonically at high-productivity firms and in early years; (ii) it is flat at mid-level firms and in middle years; and (iii) it often decreases at low-level firms and in late years. Thus our structural competition-innovation curve exhibits a "plateau" shape instead of the famous "inverted-U." Moreover, this systematic heterogeneity suggests (high) continuation values are a key factor in sustaining the (positive) competition-innovation relationship. Third, mergers become more attractive as the industry matures, and all kinds of pairs can merge. But high-type firms tend to acquire more often

[7] Computationally, the calculation of the likelihood function is the heaviest part because, for each candidate vector of parameter values, we use backward induction to solve a *nonstationary* dynamic game with 8 different types of firms and 77,520 industry states in each of the 360 periods. We perform this subroutine in C++ and the estimation procedure takes less than a week.

and low-type firms are more popular as targets, because high-type firms can gain more from increased concentration than others and low-type firms are a cheaper means for stochastic productivity gains (i.e., synergies).

In Section 5, we conduct counterfactual policy simulations to answer our main question: how far the industry should be allowed to consolidate. In Section 5.1, we find the optimal static (or "commitment") policy is to block mergers if five or less firms exist. In Section 5.2, we clarify the underlying mechanism behind this finding by decomposing the dynamic welfare tradeoff between the ex-post pro-competitive benefits of blocking mergers and the ex-ante value-destruction side effects, which reduce both competition and innovation in early years. In Section 5.3, we find that this optimal policy threshold ($\underline{N} = 5$) can be relaxed slightly if the industry is declining fast and firms are failing anyway. In Section 5.4, we find the optimal dynamic (or ex-post "surprise") policy is to initially promise no merger enforcement at all ($\underline{N}^{ante} = 1$) and then block mergers once the industry reaches 3 firms ($\underline{N}^{post} = 3$). This policy, however, relies entirely on the authority's ability to surprise (and the naïveness of firms), hence its feasibility and desirability are dubious in the long run.

We conclude in Section 6 by discussing other policy implications and limitations. The current de-facto policy of $\underline{N} = 3$ is somewhat stricter than our optimal threshold of $\underline{N} = 5$ for the HDD industry, but the welfare outcomes under these two policies are not drastically different. By contrast, our results suggest allowing mergers to duopoly or monopoly ($\underline{N} = 2$ or 1) will have a negative welfare impact that is orders of magnitude larger. Thus our main message is, "2 are few and 6 are many."[8]

1.1 Literature Context

Dynamic welfare tradeoff is a classical theme in the literature on market structure and innovation (c.f., Scotchmer 2004). Tirole (1988, p. 390) summarizes Schumpeter's (1942) basic argument that "if one wants to induce firms to undertake R&D one must accept the creation of monopolies as a necessary evil." He then proceeds to discuss this "dilemma of the patent system" but concludes that "the welfare analysis is relatively complex, and more work is necessary before clear and applicable conclusions will be within reach," (p. 399) which is exactly the purpose of this paper.

Traditional oligopoly theory suggests the main purpose of mergers is to kill competition and increase market power. Stigler (1950) added a twist to this thesis by conjecturing that, because a merger increases concentration at the industry level and non-merging parties

[8] We do not model collusion, but our finding resonates with Selten's (1973) "4 are few and 6 are many."

can free-ride on merging parties' efforts, no firms would want to take initiatives to merge. Salant, Switzer, and Reynolds (1983) proved this idea in a symmetric Cournot model, although Perry and Porter (1985) and Deneckere and Davidson (1986) revealed the fragility of the free-riding result, which crucially relied on symmetry across firms. Farrell and Shapiro (1990) used a Cournot model with cost-heterogeneity across firms, and formalized the notion of "synergy" as an improvement in the marginal cost of merging firms (above and beyond the convergence of the two parties' pre-merger productivity levels). We follow their modeling approach and definition of synergy. The latest reincarnations of this strand is Mermelstein, Nocke, Satterthwaite, and Whinston's (2014, henceforth MNSW) numerical theory of dynamic duopoly with mergers and investments, which Marshall and Parra (2015) extend to more general market structures. We provide a structural empirical companion to this literature.

Rust (1987) pioneered the empirical methods for dynamic structural models by combining dynamic programming and discrete-choice modeling, and proposed a full-solution estimation approach. Much of the empirical dynamic games literature has evolved within Ericson and Pakes's (1995) framework, and two-step methods have been developed to estimate this class of models.[9] However, typical empirical contexts of innovative industries (i.e., nonstationarity and global concentration) pose practical challenges to these methods, which led us to propose the pairing of a random-mover dynamic game (in a nonstationary environment and a finite horizon, as in Pakes 1986) with Rust's estimation approach.[10]

Applications of dynamic games to mergers include Gowrisankaran's (1995, 1999) pioneering computational work, Stahl (2011), and Jeziorski (2014). Applications to innovation include Benkard (2004), Goettler and Gordon (2011), Kim (2015), and Igami (2015, 2016). Applications to entry and exit are the largest literature, including Ryan (2012), Collard-Wexler (2013), Takahashi (2015), Arcidiacono, Bayer, Blevins, and Ellickson (2015), and Igami and Yang (2016). We have not found any empirical application of stochastically alternating-move games, but Iskhakov, Rust, and Schjerning (2014, 2016) numerically study Bertrand duopoly with "leap-frogging" process innovations. Finally, empirical evidence on patents and mergers include Ozcan (2015) and Entezarkheir and Moshiri (2015).

[9] Aguirregabiria and Mira (2007); Bajari, Benkard, and Levin (2007); Pakes, Ostrovsky, and Berry (2007); Pesendorfer and Schmidt-Dengler (2008).

[10] Egesdal, Lai, and Su's (2015) proposes MPEC algorithm as an alternative to NFXP, which is conceptually feasible but currently impractical for nonstationary, sequential-move games, due to extensive use of memory. See Iskhakov, Lee, Rust, Schjerning, and Seo (2016) for a recent tune-up to NFXP.

2 Model

This section describes our empirical model. Our goal is to incorporate a dynamic oligopoly game of mergers and innovation within Rust's (1987) dynamic discrete-choice model.

2.1 Setup

Time is discrete with a finite horizon, $t = 0, 1, 2, \ldots, T$, where the final period T is the time at which the demand for HDDs becomes zero. Each of the finite number of incumbent firms, $i = 1, 2, \ldots n_t$, has its own productivity on a discretized grid with unit interval, $\omega_{it} \in \{\omega^1, \omega^2, ...\omega^{\max}\}$, which represents the level of tacit knowledge embodied by its team of R&D engineers and manufacturing engineers. Given the productivity profile, $\omega_t \equiv \{\omega_{it}\}_{i=1}^{n_t}$, these incumbents participate in the HDD spot market and earn period profits, $\pi_{it}(\omega_t)$. Thus ω_t constitutes the payoff-relevant state variable along with the time period t, which subsumes both the time-varying demand situation and the industry-wide technological trend (i.e., Kryder's Law). We specify and estimate $\pi_{it}(\omega_t)$ in Section 4.

We assume a potential entrant (denoted by $i = 0$ and state ω^0) exists in every period and chooses whether to enter or wait when its turn-to-move arrives.[11] Upon entry, it becomes active at the lowest productivity level, $\omega_{i,t+1} = \omega^1$. If it stays out, $\omega_{i,t+1} = \omega^0$. Each of the two actions entails a sunk cost and an idiosyncratic cost shock, κ^{a0} and $\varepsilon(a_{it}^0)$, where $a^0 \in A^0 = \{enter, out\}$. An incumbent chooses between exit, innovation, merger, and staying alone without taking any major action (which we call "idling"), when its turn arrives. Each of these dynamic actions, $a \in A = \left\{exit, innovate, \{\text{propose merger to rival } j\}_{j \neq i}, idling\right\}$, entails a sunk cost, κ^a, and an idiosyncratic cost shock, $\varepsilon(a_{it})$. We follow Rust (1987) to assume $\varepsilon(a_{it}^0)$ and $\varepsilon(a_{it})$ are independently and identically distributed (i.i.d.) type-1 extreme value.

The three actions by incumbents induce the following transitions of ω_{it}. First, all exits are final and implies liquidation, after which the exiter reaches an absorbing state, $\omega_{i,t+1} = \omega^{00}$ ("dead"). Second, innovation in the HDD context means the costly implementation of re-tooling or upgrading of manufacturing equipment to improve quality-adjusted productivity,[12]

[11] In our data, entry had all but ceased by January 1996 (i.e., the beginning of our sample period) and our main focus is on the process of consolidation, but we incorporate entry to keep our model sufficiently general, so that it can be applied to the entire life cycle of an industry in principle. Another reason is that at least one episode of entry actually existed. Finis Conner founded Conner Technology in the late 1990s.

[12] We say "quality-adjusted" productivity because the industry-wide technological trend is always improving product quality (in terms of areal density) at a deterministic rate according to Kryder's Law, hence the ω_{it}s here should be understood as the de-trended version of raw productivity.

$\omega_{i,t+1} = \omega_{it} + 1$. Third, an incumbent may propose merger to one of the other incumbents and enter a bilateral bargaining. We consider two bargaining protocols: (i) Nash bargaining with equal bargaining powers between the acquirer and the target (henceforth "NB"), and (ii) take-it-or-leave-it offer by the acquirer to the target ("TIOLI").

Horizontal mergers and synergies in the HDD context are not so much about the reallocation of tangible assets (e.g., physical production capacities), which are "perishable" and tend to become obsolete within a few quarters anyway, as about combining teams of engineers who embody tacit knowledge.[13] Thus a natural way to model the evolution of post-merger productivity is to follow Farrell and Shapiro (1990) and specify $\omega_{i,t+1} = \max\{\omega_{it}, \omega_{jt}\} + \Delta_{i,t+1}$, where i and j are the identities of the acquirer and the target, respectively, and $\Delta_{i,t+1}$ is the realization of stochastic improvement in productivity. The first term on the right-hand side reflects the convergence of the merging parties' productivity levels, which Farrell and Shapiro called "rationalization," and the second term represents what they called "synergies." Given the discrete grid of ω_{it}'s (and the fact that mergers in a concentrated industry are rare events by definition), a simple discrete probability distribution is desirable, hence we specify $\Delta_{i,t+1} \sim Poisson(\lambda)$ i.i.d., where λ is the expected value of synergy.

2.2 Timing

Standard empirical models of strategic industry dynamics such as Ericson and Pakes (1995) assume simultaneous moves in each period. However, if any of the n firms can propose merger to any other firm in the same period, every proposal becomes a function of the other $n(n-1)-1$ proposals, which will lead to multiple equilibria. Instead, we consider an alternating-move game, in which the time interval is relatively short and only (up to) one firm has an opportunity to make a dynamic discrete choice within a period. Gowrisankaran (1995, 1999) and Igami (2015, 2016) are examples of such formulation with deterministic orders of moves, but researchers usually do not have theoretical or empirical reason to favor one specific order over the others. A deterministic order is particularly undesirable for analyzing endogenous mergers because early-mover advantages will translate into stronger bargaining powers, tilting the playing field and equilibrium outcomes in favor of certain firms.

For these reasons, we use stochastically alternating moves and model the timeline within each period as follows.

[13] According to Currie Munce of HGST, a big rationale for consolidation is that "As further improvement becomes technically more challenging, the industry has to pool people and talents, which would lead to further break-through." (February 27, 2015).

1. Nature chooses at most one firm (say i) with "recognition" probability, ρ_i, at the beginning of each period.

2. Mover i observes the current industry state, ω_t, forms rational expectations about its future evolution, $\{\omega_\tau\}_{\tau=t+1}^{T}$, and draws idiosyncratic shocks, $\varepsilon(a_{it})$, which represent random costs associated with the dynamic actions.[14]

3. Based on these pieces of information and their implications, mover i makes the discrete choice $a_{it} \in A_{it}$, immediately incurring the associated sunk cost, κ^a, and the idiosyncratic cost shock, $\varepsilon(a_{it})$. If i is an incumbent and chooses to negotiate a potential merger with incumbent j, the two parties bargain over the acquisition price, p_{ij}, which is a dollar amount to be transferred from i to j upon agreement. Our baseline specification of the bargaining protocol is NB, but we will also consider an alternative specification, TIOLI.[15] If the negotiation breaks down, no transfer takes place, i's turn ends without any other action or other merger negotiation, and j will remain independent.

4. All incumbent firms (regardless of the stochastic turn to move) participate in the spot-market competition, earn period profits, $\pi_{it}(\omega_t)$, and pay the fixed cost of operation, κ^c, which includes the costs of continual efforts to keep up with the industry-wide technological progress (i.e., Kryder's Law).

5. Mover i implements its dynamic action and its state evolves accordingly. If i is merging, it draws stochastic synergy, $\Delta_{i,t+1}$, which determines the merged entity's productivity in the next period, $\omega_{i,t+1}$.

These steps are repeated T times until the industry comes to an end.

2.3 Dynamic Optimization and Equilibrium

Whenever its turn to move arrives, a firm makes a discrete choice to maximize its expected net present value. Its strategy, σ_i, consists of a mapping from its effective state (a vector of

[14] Simultaneous-move entry games such as Seim (2006) characterize ε_{it}'s as private information of i, and we follow this convention. In our sequential-move context, however, ε_{it}'s do not have to be private because the other firms cannot act on ε_{it}'s even if they observed their realizations.

[15] No systematic record exists on the actual merger negotiations, and the details are likely to be highly idiosyncratic. In the absence of solid evidence, we prefer keeping the specification as neutral as possible. See Collard-Wexler, Gowrisankaran, and Lee (2014) for a non-cooperative foundation of NB.

the productivity profile ω_t, time t, and the draws of $\varepsilon_{it} = \{\varepsilon(a_{it})\}_{a \in A}$) to a choice $a_{it} \in A_{it}$—a complete set of such mappings across all t, to be precise. We may integrate out ε_{it} and consider σ_i as a collection of the ex-ante optimal choice probabilities conditional on $(\omega_{it}, \omega_{-it}, t)$.

The following Bellman equations characterize an incumbent firm's dynamic optimization problem.[16] Mover i's value *after* drawing ε_{it} is

$$V_{it}(\omega_t, \varepsilon_{it}) = \pi_i(\omega_t) + \max\left\{\bar{V}_{it}^x(\omega_t, \varepsilon_{it}^x), \bar{V}_{it}^c(\omega_t, \varepsilon_{it}^c), \bar{V}_i^i(\omega_t, \varepsilon_{it}^i), \left\{\bar{V}_{ijt}^m(\omega_t, \varepsilon_{ijt}^m)\right\}_j\right\}, \quad (1)$$

where $\bar{V}_{it}^a$s represent conditional (or "alternative-specific") values of exiting, idling, innovating, and proposing merger to rival j, respectively,

$$\begin{aligned}
\bar{V}_i^x(\omega_t, \varepsilon_{it}^x) &= -\kappa^x + \varepsilon_{it}^x + \beta E\left[\Lambda_{i,t+1}(\omega_{t+1}) | \omega_t, a_{it} = exit\right], & (2)\\
\bar{V}_i^c(\omega_t, \varepsilon_{it}^c) &= -\kappa^c + \varepsilon_{it}^c + \beta E\left[\Lambda_{i,t+1}(\omega_{t+1}) | \omega_t, a_{it} = stay\right], & (3)\\
\bar{V}_i^i(\omega_t, \varepsilon_{it}^i) &= -\kappa^c - \kappa^i + \varepsilon_{it}^i + \beta E\left[\Lambda_{i,t+1}(\omega_{t+1}) | \omega_t, a_{it} = invest\right], \text{ and} & (4)\\
\bar{V}_{ij}^m(\omega_t, \varepsilon_{ijt}^m) &= -\kappa^c - \kappa^m + \varepsilon_{ijt}^m - p_{ij}(\omega_t) + \beta E\left[\Lambda_{i,t+1}(\omega_{t+1}) | \omega_t, a_{it} = merge\ j\right]. & (5)
\end{aligned}$$

Mover i's value *before* drawing ε_{it} is

$$\begin{aligned}
EV_{it}(s_t) &= E_\varepsilon\left[V_{it}(s_t, \varepsilon_{it})\right] & (6)\\
&= \pi_i(s_t) + \gamma + \ln\left[\exp\left(\tilde{V}_{it}^x\right) + \exp\left(\tilde{V}_{it}^c\right) + \exp\left(\tilde{V}_{it}^i\right) + \sum_{j \neq i}\exp\left(\tilde{V}_{ijt}^m\right)\right],
\end{aligned}$$

where γ is Euler's constant and $\tilde{V}_{it}^a$ is the deterministic part of $\bar{V}_i^a(\omega_t, \varepsilon_{it}^a)$, that is, $\tilde{V}_{it}^a \equiv \bar{V}_i^a(\omega_t, \varepsilon_{it}^a) - \varepsilon_{it}^a$. In equations 2 through 5, $\Lambda_{i,t+1}$ represents i's expected value at $t+1$ *before* nature picks a mover at $t+1$,

$$\Lambda_{i,t+1}(\omega_{t+1}) = \rho_i(\omega_{t+1}) EV_{i,t+1}(\omega_{t+1}) + \sum_{j \neq i}\rho_j(\omega_{t+1}) W_{i,t+1}^j(\omega_{t+1}). \quad (7)$$

This "umbrella" value is a recognition probability-weighted average of mover's value (EV_{it}) and non-mover's value $\left(W_{it}^j\right)$. Nobody knows exactly who will become the mover before nature picks one. When nature picks $j \neq i$, non-mover i's value (before j draws ε_{jt} and takes

[16] Appendix B features the corresponding expressions for the potential entrant.

an action) is

$$\begin{aligned} W_{it}^{j}(\omega_t) &= \pi_i(\omega_t) - \kappa^c + E_{it}\left[\Pr\left(a_{jt} = exit\right)\right]\beta E\left[\Lambda_{i,t+1}(\omega_{t+1})\,|\omega_t, a_{jt} = exit\right] \\ &\quad + E_{it}\left[\Pr\left(a_{jt} = idle\right)\right]\beta E\left[\Lambda_{i,t+1}(\omega_{t+1})\,|\omega_t, a_{jt} = idle\right] \\ &\quad + E_{it}\left[\Pr\left(a_{jt} = invest\right)\right]\beta E\left[\Lambda_{i,t+1}(\omega_{t+1})\,|\omega_t, a_{jt} = invest\right] \\ &\quad + E_{it}\left[\Pr\left(a_{jt} = merge\ i\right)\right]p_{ji}(\omega_t) \\ &\quad + \sum_{k\neq i,j} E_{it}\left[\Pr\left(a_{jt} = merge\ k\right)\right]\beta E\left[\Lambda_{i,t+1}(\omega_{t+1})\,|\omega_t, a_{jt} = merge\ k\right], \end{aligned} \tag{8}$$

where $E_{it}\left[\Pr\left(a_{jt} = action\right)\right]$ is non-mover i's belief over mover j's choice. These value functions entail the following ex-ante optimal choice probabilities

$$\Pr\left(a_{it} = action\right) = \frac{\exp\left(\tilde{V}_{it}^{action}\right)}{\exp\left(\tilde{V}_{it}^{x}\right) + \exp\left(\tilde{V}_{it}^{c}\right) + \exp\left(\tilde{V}_{it}^{i}\right) + \sum_{j\neq i}\exp\left(\tilde{V}_{ijt}^{m}\right)}. \tag{9}$$

In equilibrium, these probabilities constitute the non-movers' beliefs over the mover's choice. We will use these optimal choice probabilities to construct a likelihood function for estimation in Section 4.3.

The bargaining protocol determines the equilibrium acquisition price, p_{ij}. Under NB, the two parties jointly maximize the following expression,

$$\begin{aligned} &\left\{\beta E\left[\Lambda_{i,t+1}(\omega_{t+1})\,|\omega_t, merge\ j\right] - p_{ij} - \beta\Lambda_{i,t+1}(\omega_{t+1} = \omega_t)\right\}^{\zeta} \\ &\times\left\{p_{ij} - \beta\Lambda_{j,t+1}(\omega_{t+1} = \omega_t)\right\}^{1-\zeta}, \end{aligned} \tag{10}$$

where $\zeta \in [0,1]$ represents the bargaining power of the acquirer (i here), which equals .5 under NB (with 50-50 split) and 1 under TIOLI. The last term in each bracket is the disagreement payoff.[17]

We solve this dynamic game for a unique sequential equilibrium in pure strategies that are type-symmetric. Note that ε_{it}'s are i.i.d. shocks whose realizations do not affect anyone's future payoff except through the actual choice a_{it}, hence we may solve this game by backward induction from the final period, T. At T, all firms' profits and continuation values are zero, so no decision problem exists. At $T-1$, a single mover (denoted by $i = T-1$) draws ε_{T-1}

[17] Only up to one deal (between i and j here) can be negotiated within a period.This setting is not as restrictive as it might seem at a first glance, because the time interval is relatively short and all other potential deals in the future are embedded in the disagreement payoff (i.e., each firm's stand-alone continuation value). This specification shares the flavor of Crawford and Yurukoglu (2012) and Ho (2009).

and takes whichever action a_{T-1} that maximizes its expected net present value. At $T-2$, another mover ($i = T-2$) draws ε_{T-2} and makes its discrete choice, in anticipation of: (i) the evolution of ω_t from $T-2$ to $T-1$, (ii) the recognition probabilities and other common factors, and (iii) the optimal CCPs of all types of potential movers at $T-1$, which imply the transition probabilities of ω_t from $T-1$ to T. This iterative process repeats itself until the initial period $t = 0$.

An equilibrium exists and is unique. First, each of the (at most) T discrete-choice problems has a unique solution given the i.i.d. draws from a continuous distribution. Second, in each period t, only (up to) one firm solves this problem in our alternating-move formulation. Third, mover t's choice completely determines the transition probability of ω_t to ω_{t+1}, but it cannot affect future movers' optimal CCPs at $t+1$ and beyond in any other way. In other words, this game is effectively a sequence of T single-agent problems. By the principle of optimality, we can solve it by backward induction for a unique equilibrium.

2.4 Other Modeling Considerations

To clarify our modeling choices, we discuss five alternative modeling possibilities that we have considered: (i) an infinite horizon, (ii) continuous time, (iii) heterogeneous recognition probabilities, (iv) alternative bargaining protocols, and (v) private information on synergies.

First, we have chosen a finite horizon over an infinite one primarily because we study the process of industry consolidation in an innovative, nonstationary industry. Another reason is multiple equilibria. Iskhakov, Rust, and Schjerning (2016) find numerous equilibria in a stochastically alternating-move duopoly game of innovation with an infinite horizon. Multiple equilibria would preclude the use of full-solution estimation methods and counterfactual analysis.

Second, continuous time modeling is an attractive alternative, but Arcidiacono, Bayer, Blevins, and Ellickson (2015) acknowledge that the feasibility of its application to a nonstationary environment is currently unknown. Another problem with a shorter or infinitesimal time interval in our context is its potential conflict with the i.i.d. idiosyncratic shocks and timing assumptions. For major and infrequent decisions such as mergers, the actual decision making and implementation take at least a month or a quarter. Shorter intervals would imply that firms draw i.i.d. random shocks every day or week. Incorporating a persistent unobserved state could alleviate this problem but create another technical challenge.

Third, some firms might be more active in M&A than others, and recognition probabilities can accommodate such heterogeneity. For example, making ρ_i depend on ω_{it} would be

conceptually straightforward, albeit computationally costly. One problem with this idea is that we have no theory. Another problem is identification. Since we have no theoretical or empirical foundation for a priori specification of asymmetric ρ_i's, we prefer keeping it symmetric and instead focus on the extent of heterogeneity *in the equilibrium CCP estimates*. Section 4.4 shows that high-type firms are more likely to acquire low-type firms, indeed.

Fourth, regarding the specifications NB and TIOLI, we may leave the bargaining powers, ζ, as free parameter and try to estimate them. However, mergers in a concentrated industry are rare events by definition, which leads to a data environment with only a handful of actual acquisition deals to estimate ζ. Thus we pre-specify NB and TIOLI as alternative models, and implement both as a sensitivity analysis.

Fifth, regarding the nature of synergies, we may consider a more complicated model of synergies with private information (e.g., some firms might privately know that their merger would yield particularly high $\Delta_{i,t+1}$), but we have chosen not to add more structures, for three reasons. First, such non-trivial private information will constitute unobserved state variables and generate a selection problem, which is an interesting problem but beyond the scope of this paper. Second, no systematic record exists on firms' subjective assessments of "chemistry," hence the identification of such factors would appear hopeless without strong additional assumptions anyway. Third, our simple model of $\Delta_{i,t+1}$ as a completely random draw actually seems the most consistent with our personal interview with Finis Conner, the co-founder of Seagate Technology, the founder of Conner Peripherals, and the founder of Conner Technology. Having founded two Fortune-500 companies in the HDD industry and engaged in some of the historical mergers, he is an embodiment of the industry's highest-quality private information. Nevertheless, he stated, "You have to dive into the water to see where the skeletons are," which means even an industry veteran would not know the internal functioning of the other firms sufficiently to predict the synergy realizations with much precision, until after the actual mergers take place.[18] Thus ours is an empirical model of Finis Conner. We keep our synergy function simple, and conduct sensitivity analysis with respect to λ in Section 4.3.

For these reasons, we believe our specification strikes the right balance amid many conceptual and practical challenges.

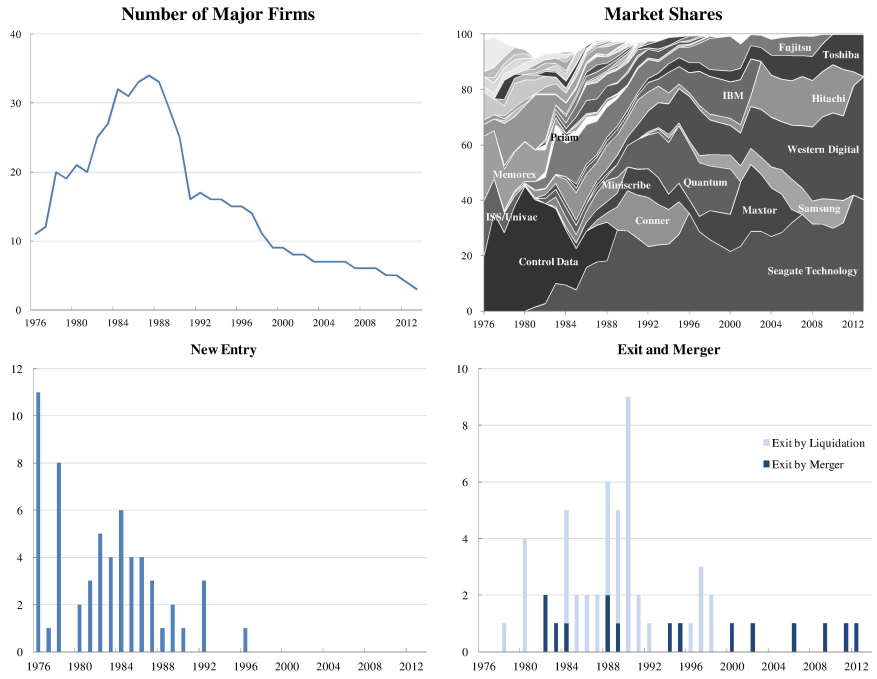
[18] From author's personal interview on April 20, 2015, in Corona del Mar, CA.

3 Data

3.1 Institutional Background and Product Characteristics

Computers are archetypical high-tech goods that store, process, and transmit data. Hard disk drives (HDDs), semiconductor chips, and network equipment perform these tasks, respectively. HDDs offer the most relevant empirical context to study mergers and innovation in the process of industry consolidation. The industry has experienced massive waves of entry and exit (liquidation), followed by consolidation among a dozen survivors since the turn of the 21st century (Figure 1).

Figure 1: Evolution of the World's HDD Industry



Note: The number of firms counts only the major firms with market shares exceeding one percent at some point of time. See Igami (2015, 2016) for the detailed analyses of product and process innovations during the 1980s and 1990s. Labels in the top-right panel indicate the names of parties to prominent merger cases.

The manufacturing of HDDs requires engineering virtuosity in fitting heads, disks, motors, and other electronic components into an air-tight black box, making them work reliably, and keeping up with the technological trend that constantly improves data-storage capacity. Kryder's Law says HDDs' areal density doubles every 12 months, which is faster than Moore's Law on the circuit density of semiconductor chips (every 18 to 24 months).

Figure 2: Product Characteristics of HDDs



Note: Left panel shows 3.5-inch HDDs of Hitachi GST, Western Digital, and Seagate Technology. Right panel shows evidence of successful marketing efforts by Microsoft and Intel (and lack thereof by HDD makers).

Despite such complexity, HDDs are also one of the simplest products in terms of economics, because they are "completely undifferentiated product," according to Peter Knight, former vice president of Conner Peripherals and Seagate Technology, and former president of Conner Technology.[19] Consumers typically do not observe or distinguish "brands" (Figure 2, left). Moreover, HDDs are physically durable but do not drive the repurchasing cycle of PCs by consumers and businesses. Microsoft and Intel ("Wintel") do, as is evident from the fact that PC users tend to be aware of the technological generations of operating systems (OS) and central processing units (CPU) but not HDDs (Figure 2, right), which means the demand for HDDs can be usefully modeled within a static framework as long as we control for the PC shipments as a demand shifter. These product characteristics inform our demand analysis in Section 4.1.

Two institutional features inform our analysis of the supply side in Section 4.2. First, unlike car manufacturing in Japan (say), the manufacturers of PCs and HDDs do not engage in long-term contract or relationship in a strict sense. The architecture of PC is highly modular and standardized interfaces connect various components including HDDs, which makes different "brands" technologically substitutable. Furthermore, "second sourcing" has long been a standard practice in the computer industry, by which a down-stream firm keeps close contact with multiple suppliers of a key component so that, in cases of accidental supply shortage from the primary seller, there will always be a backup supplier or two. According to Peter Knight, "Compaq, HP, nobody cares who makes their disk drives. They bought the lowest-price product that had reasonable quality. There was no reason for single-sourcing."[20] Second, PC makers might appear to have consolidated just as much as HDD makers, but the actual market structure of the global PC industry is more fragmented. The average

[19] From author's personal interview on June 30, 2015, in Cupertino, CA. See also Section 4.1.
[20] From author's personal interview on June 30, 2015, in Cupertino, CA. See also Section 4.1.

combined market share of the top four vendors (i.e., CR4) between 2006 and 2015 is 52.5%, which is considered between "low" and "medium" concentration. By contrast, the HDD industry's average CR4 is 91.6% during the same period.[21]

Finally, our data include some kind of solid-state drives (SSDs), but we do not explicitly model them because: (i) pure SSDs comprised less than 10% of industry sales even in the last five years of our sample period; (ii) they are made of NAND flash memory (a type of semiconductor devices), whose underlying technology is totally different from HDD's magnetic recording technology; and (iii) NAND flash memories are supplied by a different set of firms (i.e., semiconductor chip makers specialized in flash memories). Modeling SSDs means modeling the semiconductor industry. However, most of "SSDs" for desktop PCs are actually "hybrid HDDs," which combine a small NAND part with HDDs. These hybrids are part of our HDD data, and their increasing presence is captured as a secular trend of quality improvement in our data analysis.[22] Thus, when we control for the industry-wide technological trend, we are incorporating Kryder's Law for HDDs as well as Moore's Law for semiconductor devices in a reduced-form manner.

3.2 Three Data Elements

Our empirical analysis will focus on the period between 1996 and 2015 for three reasons. First, most of the exits prior to the mid-1990s were shakeout of fringe firms that occurred through plain liquidation, whereas our main interest concerns mergers in the final phase of industry consolidation. Second, the de-facto standardization of both product design and manufacturing processes had mostly finished by 1996. Specifically, the 3.5-inch form factor had come to dominate the desktop market (see Igami 2015), and manufacturing operations in South-East Asia had achieved the most competitive cost-quality balance (see Igami 2016). Third, our main data source, *TRENDFOCUS*, an industry publication series, started most of its systematic data collection at quarterly frequency in 1996.[23]

Table 1 summarizes our main dataset, which consists of three elements corresponding to three steps of our empirical analysis in the next section. Panel A is the aggregate quarterly

[21] Modeling the entire supply chain of PCs and HDDs as bilateral oligopoly would be an interesting exercise, but it is beyond the scope of this paper, whose main focus is on endogenous horizontal mergers and the long-term dynamics of innovation and entry/exit.

[22] Pure SSDs have become common for note PCs, but we focus on HDDs (including hybrids) for desktop PCs, which is still the mainstream market for HDDs.

[23] By contrast, Igami (2015, 2016) used *Disk/Trend* Reports (1977–1999), its predecessor, which published only once a year.

Table 1: Summary Statistics

Variable	Unit of measurement	Number of observations	Mean	Standard deviation	Minimum	Maximum
Panel A: Data for Demand Estimation						
HDD shipments, Q_t	Exabytes	78	15.882	17.368	0.021	53.196
HDD price, P_t	\$/Gigabytes	78	14.991	37.305	0.032	178.617
Disk price, Z_t	\$/Gigabytes	78	1.952	5.252	0.005	23.508
PC shipments, X_t	Million units	78	29.286	7.188	14.468	40.312
Panel B: Data for Variable-cost Estimation						
Market share, ms_{it}	%	590	13.2	11.0	0.0	45.7
Panel C: Data for Sunk-cost Estimation						
Indicator$\{a_{it} = merge\}$	0 or 1	1,766	0.0034	0.0582	0	1
Indicator$\{a_{it} = invest\}$	0 or 1	1,766	0.0142	0.1181	0	1
Indicator$\{a_{it} = exit\}$	0 or 1	1,766	0.0028	0.0531	0	1
Indicator$\{a_{it} = enter\}$	0 or 1	233	0.0043	0.0654	0	1
Variable profit, π_{it}	Million \$	see note	42.70	91.19	0.00	9725.22

Note: Panel A is recorded in quarterly frequency at aggregate level, Panel B is quarterly at firm-level, and Panel C is monthly at firm level. π_{it} is our period-profit estimates and contains 42,325,920 values across 7 productivity levels, 78 quarters, and 77520 industry states. Sections 4.1 and 4.2 explain how we use Panel A and Panel B to obtain π_{it}.

Source: *TRENDFOCUS* Reports (1996–2015).

data on HDD shipments, HDD price, disk price, and PC shipments,[24] which we will use to estimate HDD demand in Section 4.1. Panel B is the firm-level market shares at quarterly frequency, a graphic version of which is displayed in Figure 1 (top right). We will use demand estimates and Panel B to infer variable cost of each firm in each period in Section 4.2. Panel C is a systematic record of firms' dynamic choices between merger, R&D investment, and entry/exit, at monthly frequency. Panel C includes some elements that are derived from the other two panels, such as the indicator of R&D investment and the equilibrium variable profits.[25] We will use these dynamic choice data and stage-game payoffs to estimate the implied sunk costs associated with these actions in Section 4.3.

4 Empirical Analysis

We flesh out our dynamic oligopoly game in Section 2 with the actual data from the HDD consolidation in Section 3, which contained three elements: (i) aggregate sales, (ii) firm-level market shares, and (iii) dynamic discrete choice. Each of these data elements is paired with

[24] Appendix C features more details on Panel A, including visual plots of these variables.

[25] Appendix D.2 explains the details of this data construction, including the conversion of time intervals and the recognition of innovation and depreciation.

a model element and an empirical method to estimate demand, variable costs, and sunk costs. Table 2 offers an overview.

Table 2: Overview of Empirical Analysis

Section	Step	Model	Data	Method
4.1	Demand	Log-linear demand	Panel A	IV regression
4.2	Variable cost	Cournot competition	Panel B	First-order condition
4.3	Sunk cost	Dynamic discrete choice	Panel C	Maximum likelihood

Note: See Section 2 for the dynamic model, and Section 3 for the three data elements.

4.1 Demand Estimation

We follow Peter Knight's characterization of HDDs as "completely undifferentiated products" (see Section 3.1). To be precise, HDDs come in a few different data-storage capacities (e.g., 1 terabytes per drive), but all firms are selling these products with "the same capacities, the same speed, and similar reliability" at any given moment, so that cost is the only dimension of competition.[26] Most consumers, including the authors, do not even know which "brand" of HDDs are installed inside their desktop PCs, and PC manufacturers typically do not let consumers choose a brand. Thus homogeneous-good demand and Cournot competition are useful characterization of the spot-market transactions.

To ensure our data format is consistent with our notion of product homogeneity, we consider units of data storage (measured in gigabytes) as undifferentiated goods. We specify a log-linear demand for raw data-storage functionality of HDDs,

$$\log Q_t = \alpha_0 + \alpha_1 \log P_t + \alpha_2 \log X_t + \varepsilon_t, \tag{11}$$

where Q_t is the world's total HDD shipments in exabytes (EB, or one billion gigabytes), P_t is the average HDD price per gigabytes (\$/GB), X_t is the PC shipments (in million units) as a demand-shifter, and ε_t represents unobserved demand shocks.

Because the equilibrium prices in the data may correlate with ε_t, we instrument P_t by Z_t, the average disk price per gigabytes (\$/GB). Disks are one of the main components of HDDs, and hence their price is an important cost-shifter for HDDs. The disks are made from substrates, which are in turn made of either aluminum or glass. The manufacturers of these key inputs are primarily in the business of processing materials, and only a small fraction

[26] From author's personal interview on June 30, 2015, in Cupertino, CA.

of their revenues come from the HDD-related products. Thus we regard Z_t as exogenous to the developments within the HDD market. In Table 3, column 1 shows OLS estimates, and column 2 shows the IV estimates with disk prices as Z_t. The estimates for price-elasticity, α_1, are within the standard errors of each other. This finding of inelastic demand (i.e., $|\alpha_1| < 1$) is rationalizable under oligopoly but creates a conceptual problem under monopoly (i.e., its profit-maximizing price would be arbitrarily high). Consequently, we will ignore the top 5% of consumers with the highest willingness-to-pay to keep monopoly price finite.

Table 3: Demand Estimates

Dependent variable:	(1)	(2)	(3)	(4)
log total EB shipped	OLS	IV-1	IV-2	IV-2 rolling
log price per GB (α_1)	$-.8549^{***}$	$-.8244^{***}$	$-.8446^{***}$	see Appendix D.1
	(.0188)	(.0225)	(.0259)	
log PC shipment (α_2)	$.8430^{***}$	1.0687^{***}	$.9198^{***}$	
	(.1488)	(.1817)	(.2180)	
Constant (α_0)	-1.6452^{***}	-2.4039^{***}	-1.9033^{***}	
	(.4994)	(.6084)	(.7320)	
Number of observations	78	78	78	
Adjusted R^2	.9971	.9971	.9972	
First stage regression				
IV for HDD price	–	*Disk price*	*Time trend*	
F-value	–	3009.80	742.14	
Adjusted R^2	–	.9889	.9469	

Note: Heteroskedasticity and autocorrelation consistent standard errors are in parentheses. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively.

Although we believe disk prices represent an exogenous cost-shifter, one might still suspect the existence of some unknown endogeneity problems because, after all, disks and HDDs are close neighbors in the supply chain of computers. To address this concern, we use a logarithm of time trend as an alternative IV in column 3. This IV relies on a classical notion of technological progress as a time trend, which is particularly natural in the HDD context. Kryder's Law dictates a secular trend in the improvement of areal density (i.e., gigabytes per square inch on the surface of disks), which mechanically translates into the reduction of materials cost per gigabytes because the same number of disks can store more information.

Yet another concern is that consumers' preferences might have changed over two decades. Casual empiricism suggests people have dramatically increased the amount of data usage in everyday life, which could alter the demand parameters over time. To investigate this matter, we use rolling estimation in which we roll through the sample of 78 quarters with a 12-quarter window, using 12 observations for estimation at a time. Detailed results do not fit the table format, hence we plot the estimated coefficients against time in Appendix D.1,

and report only their time-averages in column 4 of Table 3. We use this last specification for the subsequent analysis.

Other concerns and modeling considerations that have been commonly raised include: (i) demand-side dynamics, such as durability of HDDs and the repurchasing cycle of PCs, (ii) supply-side dynamics, such as long-term contract with PC makers, and (iii) non-HDD technological dynamics, including SSDs and the semiconductor industry. Our summary views are as follows: (i) the physical durability of HDDs does not determine the dynamics of PC demand; (ii) the actual interaction between HDD makers and PC makers is more adequately described as spot-market competition rather than long-term relationship; and (iii) the non-HDD technological trend and the growing presence of hybrid HDDs are incorporated as part of Kryder's Law in our analysis. We provide further institutional and technological details on these topics in Section 3.1 (Institutional Background and Product Characteristics).

4.2 Variable Costs and Spot-market Competition

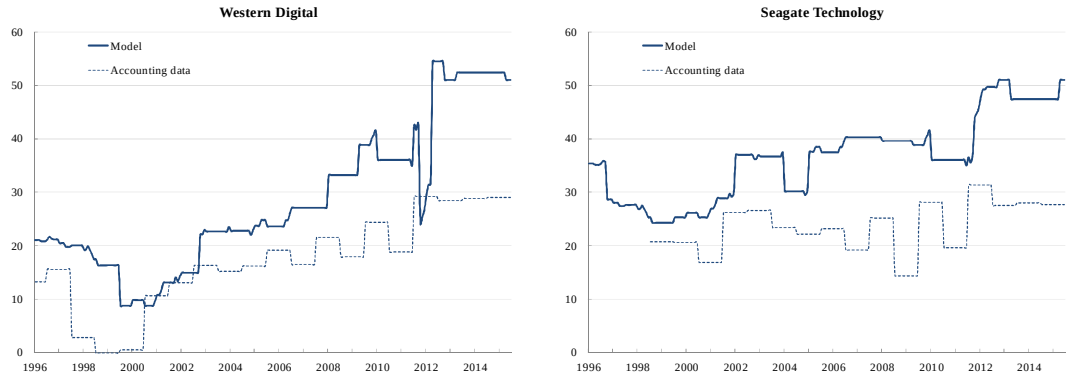
The second data element is the panel of firm-level market shares (see Figure **??**), which we now interpret through the analytical lens of Cournot competition, for two reasons. Despite selling fairly undifferentiated high-tech commodities, HDD makers' financial statements reveal positive profit margins (see dotted lines in Figure 3), which suggests Cournot model as a reasonable approximation to analyze their spot-market interactions. Another appeal of this modeling choice is that the classical oligopoly theory of mergers has mostly focused on Cournot model (see Section 1.1), which brings conceptual clarity when we follow Farrell and Shapiro (1990) to interpret the post-merger improvement in productivity as synergies, for example.

Each of the N_t firms observes the profile of marginal costs $\{mc_{it}\}_{i=1}^{N_t}$ as well as the concurrent HDD demand, and chooses production/shipping quantity, q_{it}, to maximize its variable profit

$$\pi_{it} = (P_t - mc_{it})\, q_{it} \tag{12}$$

with respect to shipping quantity q_{it}, where P_t is the price of a representative HDD and mc_{it} is the marginal cost, which we assume is predetermined and constant with respect to

Figure 3: Comparison of Profit Margins (%) in the Model and Financial Statements



Note: The model predicts economic variable profits, whereas the financial statements report accounting profits (gross profits), and hence they are conceptually not comparable. The correlation coefficient between the model and the accounting data is .8398 for Western Digital, and .5407 for Seagate Technology. With a management buy-out in 2000, Seagate Technology was a private company until 2002, when it re-entered the public market. These events caused discontinuity in the financial record.

quantity.[27] Firm i's first-order condition is

$$P_t + \frac{\partial P}{\partial Q} q_{it} = mc_{it}. \tag{13}$$

which provides one-to-one mapping between q_{it} (observed) and mc_{it} (implied) given P_t in the data and $\partial P/\partial Q$ from the demand estimates. Intuitively, the higher the firm's observed market share, the lower its implied marginal cost.

The interpretation of mc_{it} requires special attention in the high-tech context. As we discussed in Section 2 regarding synergies, "productivity" (or its spot-market equivalent, mc_{it}) in HDD manufacturing is not so much about tangible assets as about tacit knowledge embodied by teams of engineers. Thus our preferred interpretation of Cournot spot-market competition is Kreps and Scheinkman's (1983) model of quantity pre-commitment followed by price competition, given the cost profile (i.e., all active firms' productivity levels) at the beginning of each period.[28]

[27] In principle, we may replace this constant marginal cost specification with other functional forms. In the high-tech context, however, marginal costs are falling every period across the industry, and the only geographical market is the Earth. Thus one cannot rely on either inter-temporal or cross-sectional variation in data to identify the form of marginal-cost curves.

[28] One might wonder whether such "pre-committed quantities" are hard-wired to physical production capacities, and whether the latter can adjust within one to three months. Our view is that, in the context of high-tech manufacturing, *effective* physical capacities are highly "perishable" because the constant improvement in the industry's basic technology (i.e., Kryder's Law) is quite fast, which makes previously installed

Appendix D.2 shows the details of our marginal-cost estimates and how we convert them into productivity levels, ω_{it}, for the subsequent dynamic analysis. Meanwhile, we focus our main-text exposition on an external validity check of our static model. Figure 3 compares the model's predictions with accounting data, in terms of profit margins at Western Digital (left) and Seagate Technology (right), respectively. Our model takes as inputs the demand estimates and the marginal-cost estimates, and predicts equilibrium outputs, prices, and hence each firm's variable-profit margin in each year,

$$m_{it}\left(\omega_t\right) = \frac{P_t\left(\omega_t\right) - mc_{it}}{P_t\left(\omega_t\right)}, \tag{14}$$

under any market structure, ω_t (i.e., the number of firms and their productivity levels). The solid lines represent such predictions of economic profit margins along the actual history of market structure, whereas the dotted lines represent "gross profit" margins (i.e., revenue minus "cost of revenues") in the firms' financial statements.

Economic profits and accounting profits do not necessarily coincide because they are different concepts, which explains the existence of gaps in the graphs. On average, (economic) variable profit margins are higher than (accounting) gross profit margins by 11.4 and 13.8 percentage points at Western Digital and Seagate Technology, respectively, presumably because the former excludes fixed costs of operation and sunk costs of investment, whereas the latter includes some elements of fixed and sunk costs.[29] Nevertheless, the correlation coefficients between the model's prediction and the accounting data is .8398 for Western Digital, and .5407 for Seagate Technology, which confirm the relevance of (the static components of) our model.

These static analyses are interesting by themselves, but merger policy will affect not only firms' spot-market behaviors but also their incentives for mergers and investments, and hence the entire history of innovation and market structure. Thus a complete welfare analysis of industry consolidation requires endogenous mergers, innovation, and entry-exit dynamics, which will be the focus of the subsequent sections.

manufacturing equipment obsolete. Thus we prefer a rather abstract phrase "quantity pre-commitment" to "capacity" because the latter could mislead the reader to imagine "durable" physical facilities.

[29] For example, manufacturing operations in East Asia accounted for $41,304$, or 80.8%, of Seagate's $50,988$ employees on average between 2003 and 2015, whose wage bills constitute the labor component of the "cost of revenues" as a matter of accounting. However, some of these employees must have spent time and efforts on technological improvements, such as the re-tooling of manufacturing equipment for new products (i.e., product innovation), as well as the diagnosis and solution of a multitude of engineering challenges to improve the cost-effectiveness of manufacturing processes (i.e., process innovation), which should be characterized as some sorts of investment and innovation as a matter of economic interpretation.

4.3 Sunk Costs and Dynamic Discrete Choice

The third data element is the panel of firms' discrete choices between mergers, innovation, entry, and exit, which we now interpret through our dynamic structure. The previous two steps provide the period profits, $\pi_{it}(\omega_t)$, in each period under all market structure and productivity profile. In other words, the "benefits" side of the firms' choice problem is already in our hands, which leaves the "cost" side as the only remaining unknown.

Table 4: List of Parameters and Key Specifications

Parameter	Notation	Empirical approach
1. Static estimates		
Demand	$\alpha_0, \alpha_1, \alpha_2$	See Section 4.1
Variable costs	mc_{it}	See Section 4.2
Period profits	$\pi_{it}(\omega_t)$	See Section 4.2
2. Dynamics (sunk costs)		
Innovation, mergers, and entry	$\kappa^i, \kappa^m, \kappa^e$	MLE (main task of Section 4.3)
Fixed cost of operation	$\kappa^c_t(\omega_{it})$	Accounting data (see Appendix D.3)
Liquidation value	κ^x	Calibrated based on industry background
3. Dynamics (transitions)		
Discount factor (annual)	$\beta = .9$	Calibrated to the literature's standard
Prob. stochastic depreciation	$\delta = .0190$	Implied by mc_{it}
Average synergy	$\lambda = 1$	Implied by mc_{it} (sensitivity analysis with 0 & 2)
4. Other key specifications		
Terminal period	$T =$ *Dec*-2025	Sensitivity analysis with *Dec*-2020 & *Dec*-2015
Bargaining power	NB: $\zeta = .5$	Sensitivity analysis with TIOLI: $\zeta = 1$

Table 4 lists all the parameters and key specifications of our model. Before engaging in the MLE of the core parameters, $(\kappa^i, \kappa^m, \kappa^e)$, we determine the values of as many parameters as possible either as by-products of the previous two steps or directly from auxiliary data, in order to maximize the transparency of estimation process and to make computation fast and reliable. First, we pin down the other two κs as follows. The fixed cost of operations and keeping up with the industry-wide technological trend, $\kappa^c_t(\omega_{it})$, comes directly from the accounting data on sales, general, and administrative (SGA) expenses and capital expenditure (CAPEX), and are allowed to vary over time as well as across firm's productivity level.[30] We set the sunk cost of exit (or liquidation value), κ^x, to zero because tangible assets quickly become obsolete and have no obvious productive uses outside the HDD industry.

Second, three parameters govern inter-temporal transitions. The discount factor is calibrated to $\beta = .9$ at an annualized rate, a standard level in the literature. We introduce the possibility of exogenous and stochastic depreciation of productivity level at the very end of

[30] See Appendix D.3 for details.

the timeline in every period, because the marginal cost estimates exhibit occasional deterioration with probability $\delta = .0190$. Likewise, the transition patterns of mc_{it} (or equivalently, ω_{it}) indicate an average post-merger improvement of approximately \$1 (measured in terms of the discretized bin, to be precise). We "estimate" this Poisson synergy parameter, λ, as the average of the realized productivity improvements,

$$\hat{\lambda}_{MLE} = \frac{1}{n_{m\&a}} \sum_{m\&a=1}^{n_{m\&a}} \Delta_{m\&a}, \tag{15}$$

across $n_{m\&a} = 6$ actual mergers during the sample period, but one might as well consider $\lambda = 1$ as our baseline "calibration." Mergers in a concentrated industry are rare events, and IO economists are generally skeptical about the extent of synergy in the real economy, hence we conduct sensitivity analysis with $\lambda = 0$ (no synergy) and $\lambda = 2$ (strong synergy) rather than trying to discuss what its "true" and "right" value should be.

Third, two aspects of our dynamic model require fine-tuning. The first such aspect is terminal conditions. Our sample period ends in 2015Q2 but the HDD industry does not, hence we need to assume something about the post-sample end game. Our baseline specification is relatively optimistic to assume the HDD demand continues to exist until the end of year 2025, with linear interpolation between June 2015 and December 2025. Our sensitivity analyses employ more pessimistic scenarios, with $T = Dec$-2015 and Dec-2020. The second aspect is bargaining protocols. Our baseline specification is Nash bargaining with equal bargaining powers between acquirer and target, $\zeta = .5$, but we also estimate the TIOLI version, $\zeta = 1$.

Having determined the baseline configuration of the model, we proceed to estimate $(\kappa^m, \kappa^i, \kappa^e)$. The outline of our MLE procedure follows Rust (1987), who constructed the likelihood of bus engine replacement as a function of Harold Zurcher's decisions whether to replace old bus engines (choice data), their mileages (observed state variable), and the sunk cost of replacement (parameter to be estimated). Just like his identification of the replacement cost relied on the variation in the mileage of bus engines (i.e., observed differences in payoff-relevant state across time), our identification of $(\kappa^i, \kappa^m, \kappa^e)$ relies on "observed" differences in estimated period profits and their dynamic counterparts (i.e., expected net present values associated with discrete alternatives). Similarly, just like his NFXP approach nested the solution of Harold Zurcher's optimal choice problem (in the "inner loop") within the calculation of the likelihood function (to be maximized in the "outer loop"), our likelihood function nests the HDD makers' optimal choice problem. Thus our overall scheme is almost

exactly like Rust's, which is the purpose of our approach in the following sense.

We differ from Rust (1987) in three respects: (i) the HDD makers' optimal choice problem take place within a dynamic game, rather than a single agent dynamic programming problem; (ii) their opportunity to make optimal choices arrive stochastically rather than deterministically; and (iii) the underlying payoffs change over time and eventually disappear. Feature (i) fundamentally complicates the estimation problem because games generally entail multiple equilibria, which would make estimates inconsistent because one cannot use model-generated choice probabilities to pin down the parameter values if a single parameter value predicts multiple optimal choice probabilities. Our solution to this problem is threefold. First, we use an alternating-move formulation to streamline the decision problems, so that only (up to) one player makes a choice in each period. Second, we ameliorate an undesirable side effect of typical sequential-move games (with a deterministic order of moves) by making the turn to move stochastic, which led us to feature (ii) in the above. Third, we exploit the high-tech context in feature (iii) to set a finite time horizon, which shuts down the possibility of strategic interactions in the distant future and enables the backward solution of the game for a unique equilibrium. In other words, we address methodological challenges stemming from feature (i) by crafting (ii) and exploiting (iii), so that the overall scheme of estimation can proceed within the NFXP framework of Rust (1987). Thus one may consider our approach either as an extension to Rust (1987) or as an illustration of a particular class of dynamic games that is amenable to NFXP.

The optimal choice probabilities of entry, exit, innovation, and mergers in Section 2.3 (i.e., equation 9) constitute the likelihood function, contribution to which by firm i at time t is

$$l_{it}\left(a_{it}|s_t;\kappa\right)=\rho_i(s_t)\prod_{action\in A_{it}(s_t)}\Pr\left(a_{it}=action\right)^{1\{a_{it}=action\}}, \tag{16}$$

where $1\{\cdot\}$ is an indicator function. The maximum likelihood estimate (MLE) is

$$\hat{\kappa}_{MLE}=\arg\max_{(\kappa^i,\kappa^m,\kappa^e)}\frac{1}{T}\frac{1}{I}\sum_t\sum_i\ln\left[l_{it}\left(a_{it}|\omega_t;\kappa^i,\kappa^m,\kappa^e\right)\right], \tag{17}$$

where T is the number of sample periods and I is the number of firms.

The realizations of turns-to-move are not always evident in the data, and hence the implementation of MLE needs to distinguish "active" periods in which some firm took an action (such as exit, merger, or entry) and altered ω_t, and "quiet" periods in which no such

proactive moves were made by any firm. Specifically, we incorporate the random turns-to-move by setting

$$\hat{\rho}_i(s_t) = \begin{cases} 1 & \text{if } a_{it} \in \{exit, merge, enter\}, \text{ and} \\ \frac{1}{n_{\max}} \Pr(a_{it} = idle, out) & \text{if } a_{it} \in \{idle, out\} \ \forall i. \end{cases} \tag{18}$$

That is, when exit, merge, or entry is recorded in the data, we may assign probability 1 to the turn-to-move of the firm that took the action, whereas in a "quiet" period, nature may have picked any one of the firms, who subsequently decided to stay alone (or stay out) and did not alter ω_t.

In Table 5, column 1 shows our baseline estimates with (i) Nash bargaining with equal bargaining powers, $\zeta = .5$, (ii) mean synergy from the data, $\lambda = 1$, and (iii) optimistic terminal condition, $T = Dec\text{-}2025$. As a sensitivity analysis, column 2 alters ζ, columns 3 and 4 alter λ, and columns 5 and 6 alter T. All the specifications lead to similar estimates that are within the 95% confidence interval of each other.

Table 5: Maximum Likelihood Estimates of the Dynamic Parameters

Specification	(1)	(2)	(3)	(4)	(5)	(6)
Bargaining (ζ):	.5 (NB)	1 (TIOLI)	.5	.5	.5	.5
Synergy (λ):	1	1	0	2	1	1
Terminal period (T):	2025	2025	2025	2025	2020	2015
κ^i	3.0365	3.0335	3.0411	3.0326	3.0353	3.0283
	[2.65, 3.48]	[2.64, 3.47]	[2.65, 3.48]	[2.64, 3.47]	[2.65, 3.48]	[2.64, 3.47]
κ^m	5.2043	5.9239	4.9135	5.4411	5.2040	5.2049
	[4.47, 6.14]	[5.19, 6.86]	[4.18, 5.85]	[4.47, 6.14]	[4.47, 6.14]	[4.47, 6.14]
κ^e	5.2069	4.8092	5.3330	5.0667	5.3416	5.4529
	[−]	[−]	[−]	[−]	[−]	[−]
Log likelihood	−320.6502	−321.0507	−320.3957	−320.8940	−320.8981	−323.8451

Note: The 95% confidence intervals are constructed from the likelihood-ratio tests.

The directions of these differences are logical and provide an intuitive understanding of identification. For example, compare κ^m in columns 1 ($\zeta = .5$) and 2 ($\zeta = 1$). The TIOLI assumption in column 2 gives greater bargaining power to acquirers, which lowers the equilibrium acquisition price, p_{ij}, and increases the net benefits of mergers, $\bar{V}^m_{ij}$ in equation 5. Therefore, *ceteris paribus*, the TIOLI model would predict higher equilibrium CCP of merger, $\tilde{P}^m_{ij}$, but *the actual choice probabilities of merger in our data, $\bar{P}^m_{ij}$, do not change regardless of modeling assumptions*, which means the TIOLI model needs to rationalize the observed $\bar{P}^m_{ij}$ by adjusting the other components of $\bar{V}^m_{ij}$ that are free parameters. The cost of merger, κ^m, is the only free parameter that can adjust exclusively for the purpose of

rationalizing $\bar{P}^m_{ij}$ in the data, and it must become larger to keep the model's prediction, $\tilde{P}^m_{ij}$, close to $\bar{P}^m_{ij}$. Likewise, columns 3 ($\lambda = 0$) and 4 ($\lambda = 2$) feature alternative levels of expected synergy from mergers, which imply lower and higher $\bar{V}^m_{ij}$ (relative to column 1's $\lambda = 1$), respectively. As a consequence, the estimates of κ^m adjust downward (column 3) and upward (column 4), respectively, to match the model, $\tilde{P}^m_{ij}$, with the data, $\bar{P}^m_{ij}$.

The specifications in columns 5 ($T = Jan\text{-}2020$) and 6 ($T = Jan\text{-}2015$) do not exhibit as much difference from the baseline as in columns 2, 3, and 4. Our sample period ends in June 2015, hence these two alternative assumptions on (the time horizon of) the end game do not affect the estimates as long as the predicted choices during the sample period are the same. Nevertheless, the estimates of R&D investment cost, κ^i, move in an intuitive direction. The shorter the time horizon, T, the lower the continuation values as well as the incremental value associated with higher productivity, $\bar{V}^i_i$, which reduces the optimal CCP of investment, $\tilde{P}^i_i$. However, the actual probability of investment in the data, $\bar{P}^i_i$, does not change, hence κ^i has to increase to rationalize the observed investments.

The entry cost, κ^e, does not carry meaningful confidence intervals because almost any high value of κ^e could rationalize the data that contain only one entry between 1996 and 2015. That is, our estimates are a lower bound. Nevertheless, the fact that we could obtain its MLE demonstrates that one entry is more informative than zero entry. "Almost" any high κ^e can rationalize the data, but κ^e cannot be arbitrarily high, because it must permit some possibility that a potential entrant with a lucky draw of ε^e_{it} (such as Finis Conner, who founded Conner Technologies in the late 1990s) would choose to enter. The lack of precise estimate of κ^e is certainly a limitation, but it does not defeat the purpose of this paper, which is to study the process of industry consolidation in a mature market that rarely experiences new entry. Moreover, a purely theoretical possibility of new entry does not constitute a relevant threat of entry for antitrust purposes. Our conclusions from policy simulations in Section 5 will not be sensitive to κ^e, because entry almost never occurs in these simulations despite the fact that our estimate of κ^e is a lower bound.

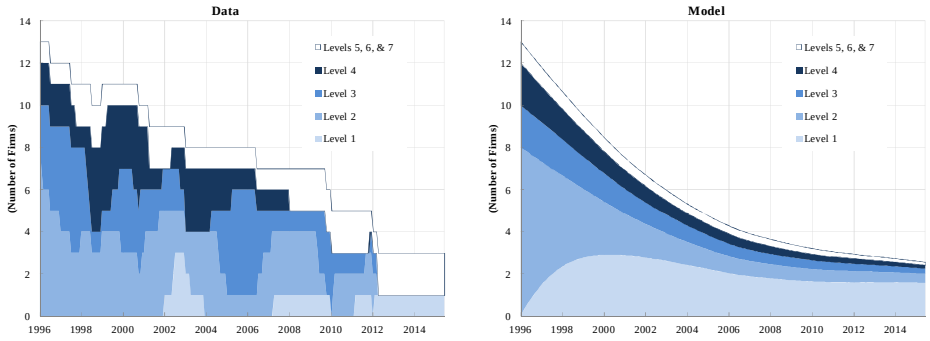
Besides these sunk costs, the NFXP estimation provides the equilibrium value and policy functions as by-products because it solves the game explicitly in the inner loop. The value function estimates contain the equilibrium enterprise values of firms in every productivity level (ω_{it}), in every period (t), under every market structure and productivity profile (ω_t). Hence, as an external validity check, we may compare these model-generated enterprise values with the actual acquisition prices in the six merger cases.[31] The comparison reveals

[31] In principle, we may use these six observed acquisition prices to estimate the bargaining parameter, ζ. However, these actual prices could reflect miscellaneous factors surrounding financial markets and are not

that at least three out of the six historical transaction values closely match the target firms' predicted values. See Appendix D.3 for further details.

Another way of assessing fit is to compare the actual and the predicted trajectories of market structure (Figure 4). The estimated model generates a smooth version of the industry consolidation process in the data, with approximately 3 firms remaining at the end of the sample period. The model also replicates some aspects of their productivity composition (e.g., the survival of a few low-level firms). We believe the estimated model provides a reasonable benchmark, with which we may compare welfare performances of hypothetical antitrust policies in Section 5.

Figure 4: Fit of the Estimated Model (Number of Firms)



Note: The model outcome is the average of 10,000 simulations based on the estimated model. The productivity types are originally defined on a discretized grid of levels 1 through 7, each step of which corresponds to an approximately $1 reduction in marginal cost.

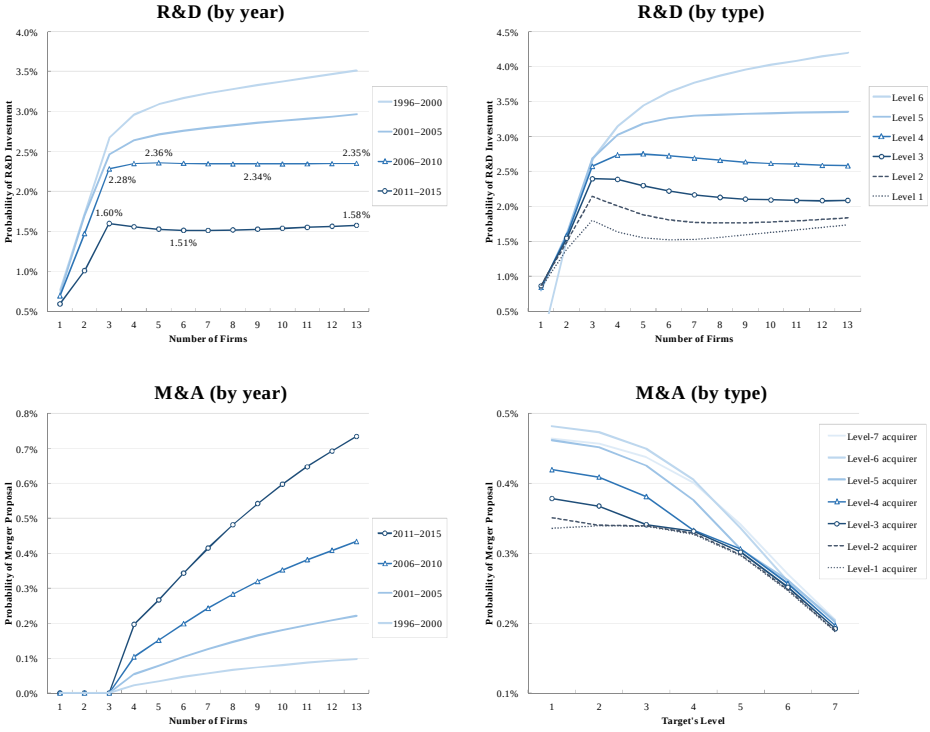
4.4 Innovation, Mergers, and Market Structure

Whereas the value-function estimates and the simulations of industry dynamics were useful for assessing fit, the policy functions are interesting on their own because they represent structural relationships between R&D incentives, M&A incentives, market structure, and the industry life cycle, which have long been a subject of both theoretical and empirical inquiries. Figure 5 shows four snapshots of the equilibrium R&D and M&A strategies by year, by type, and by market structure.

The top-left panel plots the optimal R&D investment strategy (averaged across firm types) as a function of time period and the number of active firms. Two patterns emerge: (i)

necessarily bounded between the outside options of the two parties, which is why we preset $\zeta = .5$ or 1.

Figure 5: Equilibrium R&D, M&A, and Competition



Note: This figure summarizes the overall patterns of R&D investment incentives. Our empirical analysis provides structural estimates of the equilibrium choice probabilities of each type (i.e., productivity level) of firms ω_{it}, in each market structure s_t, in each period t. Each graph pools these equilibrium investment probabilities across ω_{it} and s_t within a 5-year period and fits a nonparametric curve with local polynomials. Thus these graphs are descriptive summaries of the structural estimates rather than the estimates themselves.

the R&D decreases over time; and (ii) R&D increases with competition in early years (1996–2005), plateaus in middle years (2006–2010), and its relationship becomes non-monotonic in later years (2011–2015). These patterns merit close examination because the literature has tried (and struggled) to determine the competition-innovation relationship. The mechanism behind pattern (i) is straight-forward. The (continuation) value of an HDD firm decreases as the market shrinks and comes closer to its end, which also reduces the *incremental* value of investing in R&D and achieving higher productivity.

Pattern (ii) is more complicated. Let us first note that a standard static oligopoly model such as Cournot predicts that investments in cost-reduction decreases with the number of firms, because more rivalry means less profits as well as less *incremental* profits from innovation.[32] This mechanism is similar to Dasgupta and Stiglitz's (1980) rent-dissipation effect. However, once we include dynamic rivalry, a countervailing force sets in. Monopoly

[32] See Igami (2016) for such a stylized model with numerical examples.

will be a reluctant innovator because of Arrow's (1962) replacement effect, whereas duopoly and triopoly entail preemptive motives for innovation (Gilbert and Newbery 1982). Thus the fact that we found positive slopes implies the significance of dynamic strategic interactions.

With this underlying mechanism in mind, our interpretation of pattern (ii) becomes straight-forward. The competition-innovation curve slopes upward in early years because firms are eagerly investing in higher productivity and brighter future. The more competition, the more racing behaviors. When the industry matures, however, the option value of investment and survival starts to decline, and any additional competitors beyond 3 or 4 firms cease to affect the R&D incentives, which is why the 2006–2010 curve exhibits a plateau after 3 firms. Finally, as the industry comes closer to end, the dynamic (preemptive) incentives become less important than before, and parts of the 2011–2015 curve show symptoms of static rent-dissipation effects between 3 and 6 firms.

These competition-innovation curves do not apply uniformly across all types of firms. The top-right panel of Figure 5 reveals striking heterogeneity. High-productivity firms' R&D incentives slope upward, mid-level firms' plateau after 3 firms, and low-level firms feature non-monotonic patterns that resemble either "inverted V" or "N" shape, which are more nuanced and structural cousins of the famous "inverted U" curve (e.g., Scherer 1967; Aghion *et al.* 2005). Notice the parallel between these top panels. The competition-innovation curve slopes upward in early years and for high-type firms (i.e., under high-value circumstances and strong dynamic incentives). By contrast, the relationship becomes non-monotonic in later years or for low-type firms (i.e., under low-value circumstances and weaker dynamic considerations). Thus the continuation values and dynamic considerations (or lack thereof) are the main factor behind the complicated relationship between competition and innovation.

The incentives for mergers are also intriguing. The bottom-left panel of Figure 5 plots the optimal M&A probability (averaged across firm types) as a function of time and competition, and exhibits two patterns: (i) the more mature the industry, the more mergers; and (ii) the more rivals, the more mergers. Mergers in our model serve two purposes. One is to kill rivals and reduce competition, and the other is to draw stochastic productivity improvements. Since the latter is similar to the functionality of R&D, the former is the defining characteristic of M&A (relative to R&D). Killing a rival becomes more attractive when the fixed cost of operations become larger and makes new entry less likely, which explains pattern (i), that is, why mergers increase over time. Pattern (ii) is simpler because more rivals mean more potential merger targets. Once we divide the CCP by N_t, the merger-competition slope (per active firm) becomes virtually flat.

Who mergers with whom? The bottom-right panel of Figure 5 answers this question by plotting the CCP of merger (by *acquiring* firm's level) against the level of *target* firm. Three patterns emerge. First, all combinations of acquirers and targets are possible, as is the case in our data.[33] Second, high-level firms engage in acquisitions more often than low-level firms, because the former can expect more incremental values from reduced competition and increased productivity than the latter. Third, low- and mid-level firms are more attractive targets than high-level ones, which seems intuitive but the mechanism behind this pattern is more subtle than it appears. Eliminating a low-level rival does not soften competition as much as eliminating a high-level rival, hence the benefits from buying a low type are limited on the competition side. However, merger with a small rival can still generate as much synergy in expectation. Low-level firms are cheaper as well, because their continuation values are lower than high types'. The combination of these three factors contributed to our third finding.

We do not model all the details of M&As, which is beyond the scope of this paper. Nevertheless, the fact that such a simple model of mergers can capture and generate rich nuances seems to suggest the fruitfulness of analyzing mergers in a concentrated industry and in an explicitly dynamic setting.

5 Optimal Policy and Dynamic Welfare Tradeoff

How far should an industry be allowed to consolidate? We are now ready to simulate welfare outcomes under hypothetical merger policies, and our short answer is five. The following subsections will clarify exactly how we reach this finding (Section 5.1), the underlying mechanism that led to this finding (Section 5.2), how the outcomes may change if the industry is declining fast (Section 5.3), and the possibility of "smarter" policies (Section 5.4).

5.1 Commitment Policy

Table 6 conveys the conclusion of this subsection, that the highest net present value of social welfare is achieved under a static (or "commitment") policy in which antitrust authorities block mergers whenever the number of HDD makers is five or less (i.e., $\underline{N} = 5$). Each of the nine columns reports the discounted sums of consumer surplus (CS), producer surplus (PS), and social welfare (SW) under a hypothetical regime with $\underline{N} \in \{1, 2, ..., 9\}$ relative to the

[33] Our data contain 2 high-low, 1 mid-mid, 2 mid-low, and 1 low-mid mergers (the order of firms is acquirer-target).

baseline model ($\underline{N} = 3$).

We set $\underline{N} = 3$ in the baseline (estimated) model based on the following evidence. The FTC reports that, in merger enforcement concerning high-tech markets between 1996 and 2011, no merger was blocked until the number of "significant competitors" reached three. Specifically, (i) none of the 5-to-4 mergers was blocked; (ii) 33% of the 4-to-3 merger proposals was blocked; and (iii) 100% of the 3-to-2 and 2-to-1 proposals was blocked.[34] Thus $\underline{N} = 3$ is an accurate description of the actual policy during our sample period. This *de facto* rule of the game is a shared perception among antitrust practitioners and firms in Silicon Valley, according to our conversations with former chief economists at the FTC and the Antitrust Division of the DOJ, antitrust economic consultants, as well as senior managers at the HDD manufacturers.

Table 6: Welfare Performance of Commitment Policies

Threshold number of firms ($\underline{N}$)	1	2	3	4	5	6	7	8	9
Discounted Jan-1996 value									
Consumer surplus (%)	▲24.12	▲7.40	±0	+0.45	+0.65	+0.73	+0.77	+0.80	+0.82
Producer surplus (%)	+118.46	+32.68	±0	▲3.11	▲5.26	▲5.55	▲8.12	▲9.19	▲10.39
Social welfare (%)	▲14.64	▲4.74	±0	+0.21	+0.26	+0.24	+0.18	+0.13	+0.07
Undiscounted sum									
Consumer surplus (%)	▲79.61	▲20.74	±0	+1.14	+1.62	+1.84	+1.97	+2.06	+2.11
Producer surplus (%)	+278.70	+64.64	±0	▲4.87	▲7.62	▲9.21	▲10.67	▲11.68	▲12.64
Social welfare (%)	▲48.16	▲13.24	±0	+0.62	+0.81	+0.87	+0.86	+0.85	+0.81

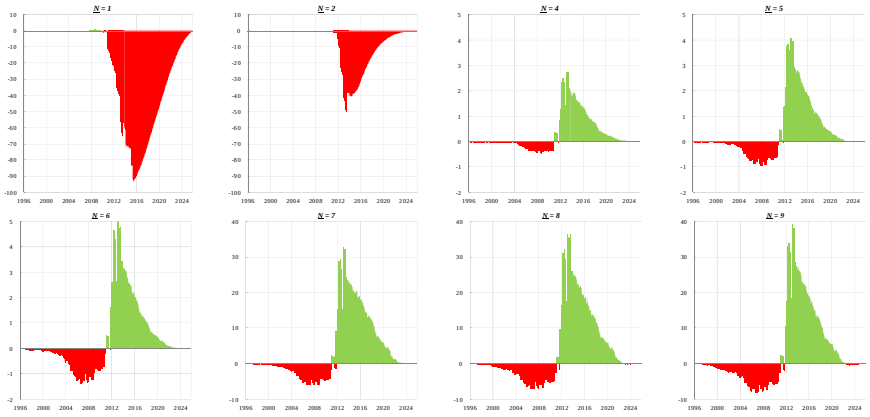
Note: All welfare numbers are expressed in terms of percentage change from the baseline outcome under $\underline{N} = 3$. These changes might appear small because most of the counterfactual policies' impacts are concentrated in later periods, at which point the market size is shrinking and discounting attenuates their values as of January 1996.

Computational implementation is straightforward. We estimated the baseline model by searching over the parameter space of $(\kappa^i, \kappa^m, \kappa^e)$ to maximize the likelihood of observing the actual choice patterns in the data (in the outer loop), and by solving the dynamic game by backward induction to calculate the predicted choice patterns based on the model (in the inner loop), in which the sunk cost of merger is κ^m when $n_t > 3$ but ∞ when $n_t \leqslant 3$. Simulating welfare outcomes under an alternative regime is simpler than estimation. First, solve the counterfactual game with the same parameter estimates $\left(\hat{\kappa}^i, \hat{\kappa}^m, \hat{\kappa}^e\right)$ but in a different policy environment ($\underline{N} \neq 3$) just once, and obtain the optimal choice probabilities in the counterfactual equilibrium. Second, use these CCPs to simulate 10,000 counterfactual

[34] See Federal Trade Commission (2013), Table 4.7 entitled "Number of Significant Competitors in Electronically-Controlled Devices and Systems Markets." Our model can incorporate similar policy regimes based on HHI thresholds instead of $\underline{N}$, but we found no clear HHI threshold in FTC (2013).

industry histories, $\{\omega_t\}_{t=0}^{T}$. Third, calculate $\{(CS_t, PS_t, SW_t)\}_{t=0}^{T}$ along each simulated history, take their average across the 10,000 simulations, and summarize their (undiscounted) time profiles in terms of time-0 discounted present values.

Figure 6: Time Profile of Undiscounted Welfare (% change)



Note: Each panel shows the average of 10,000 simulations under an alternative policy regime.

The threshold policy $\underline{N} = 5$ optimally trades off the ex-post pro-competitive effect of blocking mergers with the negative ex-ante value-destruction side effects. Figure 6 visualizes the dynamic welfare tradeoff by showing the time profiles of undiscounted social welfare, again relative to the baseline performance under $\underline{N} = 3$. The most staggering patterns emerge from the more permissive regimes ($\underline{N} = 1$ and 2), in which the deadweight losses under monopoly and duopoly during the second half of the sample period dominates any positive changes during the first half. Thus the first and the most obvious finding from our welfare analysis is that allowing mergers to monopoly or duopoly is a bad idea, even if we account for potentially positive side effects on ex-ante incentives to enter and innovate.

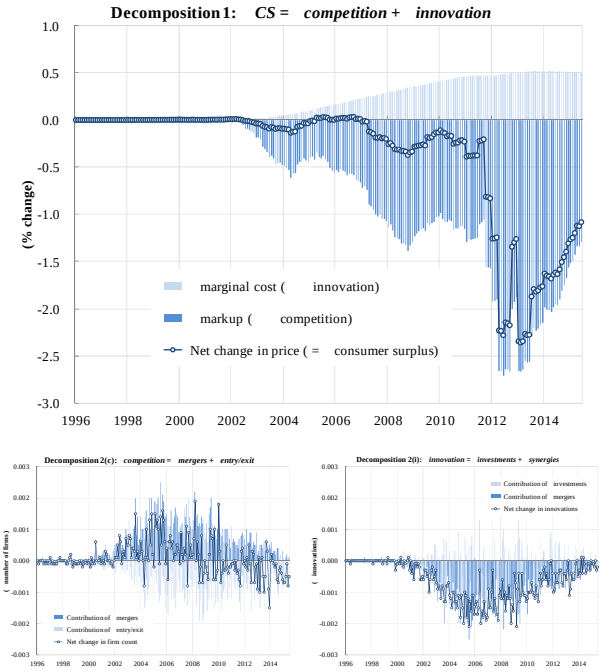
By contrast, stricter policies ($\underline{N} = 4, 5$, and 6) generate more nuanced and qualitatively interesting welfare tradeoffs between positive ex-post performances and negative ex-ante side effects. The positive changes reflect the direct impact of blocking mergers after the number of firms comes down to 4, 5, or 6. But this pro-competitive effect of antitrust policy is partially offset by preceding negative changes, which seem to suggest some side effects that discourage entry and R&D investments early on. Both the positive and the negative effects become larger as the policy threshold rises from four to six, but their rates of change are not always in balance. That is, the negative side effects grow faster than the positive main effect, to the extent that the net impact on discounted social welfare peaks at $\underline{N} = 5$ and then

declines at higher thresholds (Table 6). Thus the second and the more nuanced finding from our welfare analysis is the optimality of $\underline{N} = 5$ and the subtle tradeoff that it epitomizes. That is, tougher merger policy is not unambiguously better. The next subsection further dissects its underlying mechanisms.

5.2 Decomposing Dynamic Welfare Tradeoff

Figure 7 illustrates this dynamic welfare tradeoff, first by decomposing the changes in consumer surplus into the changes in competition and innovation, and then by further decomposing: (i) the changes in competition into the changes in mergers and entry/exit, and (ii) the changes in innovation into the changes in mergers and in-house investments.

Figure 7: Decomposition of Dynamic Welfare Tradeoff



Note: We use the (un-weighted) average marginal cost across firms. Alternative summary statistics such as the minimum or market share-weighted average do not qualitatively alter the decomposition patterns.

The summary of welfare performances under alternative policy regimes in the previous subsection indicated the presence of complicated dynamic welfare tradeoff. This subsection clarifies its underlying mechanisms by focusing on changes in CS, a good "summary statistic"[35] of which is price (i.e., the average HDD price per gigabytes) in the current em-

[35] We do not intend to claim that Δp completely determines ΔCS. It does not, because we allow both the

pirical context with homogeneous goods (i.e., bytes of data-storage functionality in personal computing).

The top panel of Figure 7 Shows the CS-performance of the optimal policy ($\underline{N} = 5$) relative to the baseline ($\underline{N} = 3$). The line graph represents the change of price (Δp_t), and the two bar graphs show its tow determinants, which are the average markup (Δm_t) and the average marginal cost (Δmc_t) across active firms.[36] Their accounting relationship permits our first decomposition,

$$\Delta CS_t \propto \Delta p_t = \Delta m_t + \Delta mc_t. \tag{19}$$

In words, CS increases when the price decreases, which can be the result of either reduced markups (i.e., increased "competition"), reduced costs (i.e., increased productivity or "innovation"), or both. Thus we are decomposing ΔCS into changes in "competition" and "innovation." The graph suggests Δm_t and Δmc_t tend to move in opposite directions and cancel each other.

The two bottom panels of Figure 7 further decompose the "competition" and the "innovation" effects of the optimal policy regime. On the "competition" side, Δm_t reflects changes in the number of firms, Δn_t, which in turn reflects the changes in *net entry* ($= entry - exit$), ΔNE_t, and mergers, ΔMA_t,

$$\Delta m_t \propto \Delta n_t = \Delta NE_t + \Delta MA_t. \tag{20}$$

In the bottom-left panel, the line graph shows Δn_t, and the two bar graphs show ΔNE_t and ΔMA_t, respectively. The contribution of the policy to competition (Δn_t) is mostly positive from the merger channel (ΔMA_t) because more mergers are blocked under $\underline{N} = 5$ than under $\underline{N} = 3$, which represents the classical pro-competitive effect of blocking mergers.

However, this gain is partially offset by the negative contribution from the entry/exit channel (ΔNE_t). The mechanism behind this countervailing effect is partial destruction of firms' continuation values. The reduced opportunities for mergers mean reduced profit margins due to more competition in the future, reduced opportunities for productivity growth by stochastic synergy draws, and reduced opportunities for profitable exit by being acquired,

covariate (i.e., X_t) and the parameters (i.e., $\hat{\alpha}_t$s) to change over time. Here we mean "summary statistic" for illustration purposes only.

[36] The average markup and cost do not reflect all of the changes in individual firms' performances, but these measures are "sufficient" for our current purpose of illustrating the key determinants of welfare under different policy regimes.

all of which reduce the option value of entry and survival. Entry and survival require lump-sum and continual investments (i.e., κ^e and κ^c) to catch up and keep up with the industry-wide technological trend, respectively. Less rosy prospect for the future profit opportunities reduces expected benefits but not expected costs, which is the reason for less entry and more exits, hence $\Delta NE_t < 0$.

The microeconomic mechanism is similar but more complicated and interesting on the "innovation" side. Δmc_t primarily reflects the productivity levels of firms, $\Delta\omega_{it}$s, which can change either as a result of in-house R&D investments, ΔRD_t, or stochastic synergy draws upon successful mergers, ΔMA_t,

$$\Delta mc_t \propto \Delta\omega_{it} = \Delta RD_t + \Delta MA_t. \tag{21}$$

The bottom-right panel of Figure 7 features a line graph representing the counts of $\Delta\omega_{it}$, and two bar graphs representing ΔRD_t and ΔMA_t, respectively. The $\underline{N} = 5$ policy's contribution to "innovation" is negative on the synergy side because it blocks more mergers than under the $\underline{N} = 3$ regime.

However, the in-house R&D channel does not seem to sufficiently offset these forgone synergies, for two reasons. First, the policy promotes more competition ex post, but this increased competition does not necessarily translate into increased R&D. As Figure 5 illustrated, the equilibrium R&D investment CCPs all but stop increasing after three or four firms, and occasionally decrease with more than five firms. Second, the aforementioned destruction of firms' continuation values re-surfaces here and discourages in-house investments. The side effects of value destruction are not limited to entry and survival; they affect all kinds of investments including in-house R&D, which is an investment in productivity growth and hence incremental profits and values. The overall level of continuation values decreases, and so does the incremental values from making such investments. Thus the impact of value-destruction manifests itself through multiple side effects, which is the reason the "optimal" policy ($\underline{N} = 5$) does not substantially outperform its neighboring thresholds such as $\underline{N} = 3$ and $\underline{N} = 4$.

5.3 Failing Firms and Declining Industries

In a mature industry such as HDDs, regulators often have to deal with "failing firms," that is, firms which: (i) are in imminent danger of failure (in a more severe condition than insolvency and close to ceasing operations), (ii) cannot be reorganized in Chapter 11

bankruptcy, and (iii) cannot find alternative purchaser (or other less anticompetitive uses) of their assets.[37] To our knowledge, no formal economic analysis exists on this subject, possibly because a systematic evaluation of failing firms requires modeling and estimation of nonstationary industry dynamics. Exits (through liquidation) in our model meets all of the three criteria for "failing firms," hence our model can handle such cases, in principle. However, the equilibrium CCPs of exit are less than 10% in almost all states and periods in our baseline estimate. Consequently, we have chosen not to study failing firms per se but to ask a broader question regarding the optimal policy toward declining industries, in which exits become more likely.[38]

Table 7: Commitment Policies in Fast-declining Industry

Threshold number of firms ($\underline{N}$)	1	2	3	4	5	6	7	8	9
$T = Dec$-2020									
Consumer surplus (%)	▲17.37	▲6.18	±0	+0.42	+0.77	+0.84	+0.89	+0.89	+0.93
Producer surplus (%)	+90.89	+28.83	±0	▲3.08	▲5.93	▲7.33	▲8.89	▲10.13	▲11.63
Social welfare (%)	▲10.28	▲3.89	±0	+0.19	+0.33	+0.31	+0.25	+0.17	+0.10
$T = Dec$-2015									
Consumer surplus (%)	▲10.58	▲4.52	±0	+0.34	+0.48	+0.49	+0.54	+0.54	+0.53
Producer surplus (%)	+64.42	+23.86	±0	▲2.77	▲4.66	▲5.93	▲7.58	▲8.87	▲10.35
Social welfare (%)	▲5.73	▲2.69	±0	+0.14	+0.15	+0.08	+0.01	▲0.07	▲0.17

Note: All welfare measures net present values as of January 1996, expressed in terms of percentage change from the baseline outcome under $\underline{N} = 3$.

Should the authority relax its merger policy when the industry is declining fast? We will answer this question in two steps. First, we capture the notion of "declining industry" (and hence higher exit rates in equilibrium or "failing firms") by hypothetically eliminating much of the HDD demand in the post-sample period (i.e., after June 2015). Our baseline model assumes the demand will linearly decline to zero between June 2015 and December 2025, reflecting what we presume to be an unspoken consensus forecast among industry participants. By contrast, this subsection simulates alternative industry dynamics in which the demand converges to zero in December 2020 and December 2015 (i.e., $T = Dec$-2020 and Dec-2015), five and ten years earlier than our baseline scenario. We solve these new games to find their respective equilibrium CCPs, simulate 10,000 histories, and calculate their average welfare performances. We maintain the baseline policy regime ($\underline{N} = 3$) throughout these procedures. Second, within each of these hypothetical demand scenarios (i.e., $T = Dec$-2020

[37] See McFarland and Nelson (2008) for legal details.

[38] "Declining industries" do not constitute a valid defense in the U.S. legal context (except under a brief period during the Great Depression), and the permission of "recession cartels" in Japan was repealed in 1999, so we are not using this phrase in a legal sense.

and *Dec*-2015), we compute welfare outcomes under alternative policy regimes (i.e., $\underline{N} \neq 3$), so that we can determine the optimal policy under each end-game scenario.

Table 7 demonstrates how the dynamic welfare tradeoff alters its balance, albeit slightly, when the industry is declining faster. The optimal static threshold continues to be five, but permitting mergers to quadropoly, triopoly, or even duopoly would not be as harmful as in Table 6, because there are not too many consumers to be harmed when the demand is disappearing precipitously. The fact that $\underline{N} = 5$ continues to be optimal might appear surprising, but the contribution of the last $5 \sim 10$ years of the industry's life cannot be too large in terms of net present value as of January 1996. Hence it is actually surprising to witness *any* qualitative change, such as the negative SW performance of $\underline{N} = 8$ and 9 under $T = Dec$-2015 scenario, which does convey the idea that the optimal policy is likely to be a more relaxed threshold rather than a stricter one.

5.4 Opportunistic Policy

Thus far we have considered only a static (or time-invariant) policy design that commits the authority to a particular merger threshold. We have intentionally kept our discussions within such static thresholds because of their simplicity and direct connection to the practitioners' rule of thumb. One should also note that a detailed analysis of dynamic welfare tradeoff becomes quite complicated even under such a simple policy design concept. Nevertheless, a sophisticated or theoretically inclined reader might be wondering if the authority can do better than simply committing to $\underline{N} = 5$, and our short answer is "yes (in the short run) and no (in the long run)." Let us explain what we mean in the following.

Finally, Table 8 considers "smart" policies in which the authority acts opportunistically and alters the merger threshold ex post.[39] The optimal surprise policy is to initially promise no antitrust scrutiny at all (i.e., declare $\underline{N}^{pre} = 1$) and then start blocking mergers when the industry consolidation reaches triopoly (i.e., $\underline{N}^{post} = 3$). In this way, the elusive quest for monopoly profits will attract massive entry and innovation early on (i.e., no value-destruction side effects), and the surprise ban on mergers after triopoly will ensure sufficient pro-competitive outcomes ex post.

Computationally, we implement these opportunistic policies as follows. First, we start

[39] We refrain from simulating more complicated policies (and their possible strategic interactions with the firms) because intuitive understanding of the results will become increasingly more difficult, the actual policy implementation will become impractical, and we could not find anecdotal or quantitative evidence that would suggest such strategic interplay. Nevertheless, these ideas do stimulate theoretical curiosity, and we refer the reader to MNSW (2014) and Jeziorski (2014) for such investigations.

Table 8: Performance of Opportunistic Policies

Promised threshold ($\underline{N}^{pre}$)	1	1	1	1	1	1	1	1	1
True threshold ($\underline{N}^{post}$)	1	2	3	4	5	6	7	8	9
Consumer surplus (%)	▲24.12	▲7.86	+1.04	+0.69	+0.75	+0.80	+0.86	+0.91	+0.91
Producer surplus (%)	+118.46	+34.50	▲5.24	▲4.52	▲5.95	▲7.11	▲8.71	▲9.96	▲11.10
Social welfare (%)	▲14.64	▲5.04	+0.62	+0.34	+0.30	+0.27	+0.23	+0.19	+0.11

Note: All welfare measures net present values as of January 1996, expressed in terms of percentage change from the baseline outcome under $\underline{N} = 3$ (both promised and true).

simulating the industry's history by using the equilibrium CCPs under $\underline{N} = \underline{N}^{pre} = 1$, which corresponds to the $\underline{N} = 1$ counterfactual in Section 5.1. Second, whenever the simulated number of firms reaches the true (but unannounced) threshold, $\underline{N}^{post}$, our algorithm switches to the equilibrium CCPs under $\underline{N} = \underline{N}^{post}$ and keeps simulating the history until $t = T$, which corresponds to the other counterfactuals with $\underline{N} > 1$ in Section 5.1. Third, collect 10,000 simulated histories and calculate their average welfare performance. This is the outcome that we attribute to each pair $\left(\underline{N}^{pre}, \underline{N}^{post}\right)$ that represents a particular ex-post policy.

To some readers, this simulation experiment might appear too complicated and unrealistic at a first glance, but negative surprises are facts of life. In the American political context, for example, consider a long spell of the Republican "pro-business" regime, followed by stronger regulatory oversight under the Democratic regime. Another example is the emergence of the Chinese antitrust policy in 2008. Its Ministry of Commerce (MOFCOM) almost stopped the latest HDD merger between Western Digital and Hitachi GST in 2012, which had already been cleared by the authorities in the U.S., Japan, South Korea, and Europe.

In the long run, however, such a "smart" policy is not going to be wise, because governments cannot fool financial markets forever. One industry might be tricked, but the subsequent cohorts of high-tech industries may not. The authority can surprise only once.

6 Conclusion

This paper proposed an empirical model of mergers and innovation to study the process of industry consolidation, with the HDD as a working example. We used quantitative methods to clarify the dynamic welfare tradeoff inherent in antitrust policy, and determined the welfare-maximizing merger threshold to be five in our empirical context. This finding is specific to the parameters of consumers' preferences, production technology, and investment technology in our data, hence each high-tech industry requires careful modeling and mea-

surement, just like the actual enforcement of antitrust policy proceeds on a case-by-case basis.

The fact that the authority has permitted mergers to triopoly in the HDD market (i.e., beyond our optimal threshold) does not appear particularly troubling, because the magnitude of welfare differences is small as long as the threshold is three or higher. Thus the danger of type II errors (i.e., not rejecting what needs to be rejected) is not overwhelming. By contrast, permitting mergers to duopoly or monopoly would lead to negative welfare impacts that are larger by an order of magnitude. Our model focuses on the "direct" effect of mergers on market structure (i.e., the number of competitors and market shares) and does not incorporate the "coordinated" effect that might facilitate collusive conducts, such as those studied by Miller and Weinberg (2015). Hence the negative effect on consumer surplus in our study represents a lower bound, and the actual harm of monopoly and duopoly could be greater.

Moore's Law (or its HDD-equivalent, Kryder's Law) is another important subject beyond the scope of this paper. Our empirical framework can incorporate such secular technological trends on a larger scale without any conceptual difficulty, because the only change will be re-defining a larger state space that can span a wider range of productivity levels. A larger state space, however, creates a computational problem. Our current implementation already uses approximately 48 gigabytes of physical memory (DRAM). A drastic expansion of state space requires more data-storage capacity with faster access speed. Consequently, our structural estimation of Moore's Law has to wait until Moore's Law makes such a venture possible.[40]

[40] We would also need new IVs for demand estimation if we endogenize Moore's Law.

Appendices: Table of Contents

Appendix A lists our personal interviews with industry veterans. Appendices B, C, and D supplement the details of Sections 2 (Model), 3 (Data), and 4 (Empirical Analysis), respectively.

Appendix A List of Interviews

For confidentiality reasons, we do not quote from our personal interviews with the industry sources. The only exceptions are historical overviews and events in the distant past (by the standard of Silicon Valley). Nevertheless, almost every modeling choice, parameterization, and estimation result has tight connections to the actual data-generating process, which we learned through these interviews.

Table 9: Interviews with Industry Sources

#	Date	Location	Name	Affiliation (position)
1	Various	TRENDFOCUS office (Cupertino, CA)	Mark Geenen John Kim John Chen Don Jeanette	TRENDFOCUS (president & VPs) Microscience International Komag Toshiba, Fujitsu
2	1/22/2015	Fibbar MaGees Irish pub (Sunnyvale, CA)	Reggie Murray	Ministor (founder) Maxtor (thin-film head) Memorex
3	2/27/2015	HGST/IBM office (San Jose, CA)	Currie Munce	HGST/IBM (SSD)
4	3/5/2015	SIEPR (Stanford, CA)	Lawrence Wu	NERA Consulting (president)
5	3/11/2015	SIEPR (Stanford, CA)	Orie Shelef	Former merger consultant
6	3/23/2015	Residence (Monte Sereno, CA)	Tu Chen	Komag (founder)
7	4/17/2015	Seagate headquarters (Cupertino, CA)	Jeff Burke	Seagate (VP of strategic marketing & research)
8	4/20/2015	Residence (Corona del Mar, CA)	Finis Conner	Conner Technology (founder) Conner Peripherals (founder) Seagate (co-founder) International Memories Inc. Shugart Associates (co-founder)
9	6/30/2015	BJ's restaurant & brewery (Cupertino, CA)	Peter Knight	Conner Technology (president) Conner Peripherals (senior VP) IBM
10	7/1/2015	Gaboja restaurant (Santa Clara, CA)	MyungChan Jeong	HGST/IBM (R&D engineer) Seagate, Maxtor Samsung Electronics

Note: Affiliations are listed from new to old. VP stands for vice president.

Appendix B Supplementary Materials for Section 2

Potential Entrant's Problem

Section 2.3 focused on the exposition of incumbent firms' problem. This section explains the detail of potential entrant's problem.

If nature picks a potential entrant i as a proposer, i draws $\varepsilon_{it}^0 = (\varepsilon_{it}^e, \varepsilon_{it}^o)$ and chooses to enter or stay out, which entail the following alternative-specific values,

$$\bar{V}_i^e(s_t, \varepsilon_{it}^e) = -\kappa^e + \varepsilon_{it}^e + \beta E\left[\Lambda_{i,t+1}(s_{t+1}) | s_t, a_{it} = enter\right], \text{ and} \tag{22}$$

$$\bar{V}_i^o(s_t, \varepsilon_{it}^o) = \varepsilon_{it}^o + \beta E\left[\Lambda_{i,t+1}(s_{t+1}) | s_t, a_{it} = out\right], \tag{23}$$

respectively. Thus the potential entrant's value *after* drawing ε_{it}^0 is

$$V_{it}^0\left(s_t, \varepsilon_{it}^0\right) = \max\left\{\bar{V}_i^e(s_t, \varepsilon_{it}^e), \bar{V}_i^o(s_t, \varepsilon_{it}^o)\right\}, \tag{24}$$

and its expected value *before* drawing ε_{it}^0 is

$$EV_{it}^0(s_t) = E_\varepsilon\left[V_{it}^0\left(s_t, \varepsilon_{it}^0\right)\right] = \gamma + \ln\left[\exp\left(\tilde{V}_{it}^e\right) + \exp\left(\tilde{V}_{it}^o\right)\right]. \tag{25}$$

These expressions correspond to equations 1 through 6 in the main text.

When the non-mover is a potential entrant, its non-mover expected value is simpler than incumbent's version in equation 8,

$$\begin{aligned} W_{it}^{0j}(s_t) = {} & \sigma_{it}(a_{jt} = exit)\beta E\left[\Lambda_{i,t+1}(s_{t+1}) | s_t, a_{jt} = exit\right] \\ & + \sigma_{it}(a_{jt} = stay)\beta E\left[\Lambda_{i,t+1}(s_{t+1}) | s_t, a_{jt} = stay\right] \\ & + \sigma_{it}(a_{jt} = invest)\beta E\left[\Lambda_{i,t+1}(s_{t+1}) | s_t, a_{jt} = invest\right] \\ & + \sum_{k \neq i,j} \sigma_{it}(a_{jt} = merge\ k)\beta E\left[\Lambda_{i,t+1}(s_{t+1}) | s_t, a_{jt} = merge\ k\right], \end{aligned} \tag{26}$$

because it does not earn profit, pay fixed cost, or become a merger target.

When nature picks a potential entrant j as a mover, equations 8 and 26 become

$$\begin{aligned} W_{it}^{j}(s_t) &= \pi_i(s_t) - \kappa^c \\ &\quad + \sigma_{it}\left(a_{jt}^0 = enter\right)\beta E\left[\Lambda_{i,t+1}(s_{t+1})\,|s_t, a_{jt}^0 = enter\right] \\ &\quad + \sigma_{it}\left(a_{jt}^0 = out\right)\beta E\left[\Lambda_{i,t+1}(s_{t+1})\,|s_t, a_{jt}^0 = out\right], \text{ and} \end{aligned} \tag{27}$$

$$\begin{aligned} W_{it}^{0j}(s_t) &= \sigma_{it}\left(a_{jt}^0 = enter\right)\beta E\left[\Lambda_{i,t+1}(s_{t+1})\,|s_t, a_{jt}^0 = enter\right] \\ &\quad + \sigma_{it}\left(a_{jt}^0 = out\right)\beta E\left[\Lambda_{i,t+1}(s_{t+1})\,|s_t, a_{jt}^0 = out\right] \end{aligned} \tag{28}$$

for an incumbent non-mover and a potential entrant non-mover, respectively.

These value functions entail the following optimal choice probabilities before potential-entrant mover i draws ε_{it}^0,

$$\Pr\left(a_{it}^0 = action\right) = \frac{\exp\left(\tilde{V}_{it}^{action}\right)}{\exp\left(\tilde{V}_{it}^{e}\right) + \exp\left(\tilde{V}_{it}^{o}\right)}, \tag{29}$$

which corresponds to equation 9.

Appendix C Supplementary Materials for Section 3

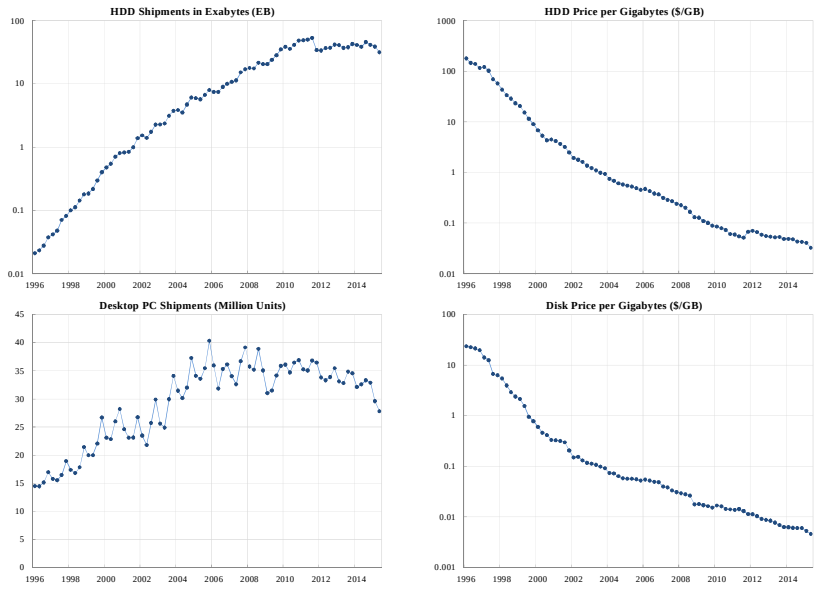
Data Patterns Underlying Demand Estimation (Panel A)

Figure 8 summarizes data patterns of Panel A, that is, the four variables for demand estimation, (Q_t, P_t, X_t, Z_t). The HDD shipment volume in EB (Q_t) has grown steadily on the back of PC shipments (X_t) as the upper- and lower-let panels show. The HDD price per GB (P_t) has been decreasing as a result of Kryder's Law. With this secular trend in storage density, the disk price per GB (Z_t) has fallen dramatically, because more data can be stored on the disk surface of the same size. The upper- and lower-right panels capture these trends. Thus the downward trends in P_t and Z_t reflect both process innovation (i.e., lower marginal costs) and product innovation (i.e., higher "quality" or data-storage capacity per HDD unit) in this industry.

Market Shares Before and After Mergers (Panel B)

In Section 3, we visualized and summarized the data patterns of firm-level market shares (Panel B) in Figure 1 and Table 1. In this section, we supplement these exhibits with the list of 14 merger cases and the Herfindahl-Hirschman Index (HHI).

Figure 8: Data for Demand Estimation at the Level of Gigabytes (GB)



Note: See Sections 3.2 and 4.1 for summary statistics and demand estimation, respectively.

Table 10 shows the combined market share of the acquiring firm and the target firm declined after merger in each of the 14 cases, which suggests the theoretical prediction of free-riding by the non-merging parties could be a real phenomenon. At the same time, the acquiring firms managed to achieve expansions relative to their individual pre-merger market shares, which is consistent with our interviews with the industry participants, in which they explained gaining market shares as the primary motivation for mergers. Finally, a larger firm acquires a smaller firm in most of the cases, which seems intuitive.

Figure 9 overlays the historical HHI on the number of firms, n_t. The HHI correlates negatively with n_t by construction. It started at around $2,000$ in the late 1970s, decreased to $1,000$ in the mid 1980s due to massive entry, and was mostly unaffected by the shakeout because fringe firms' liquidation-exit did not really change the surviving firms' market shares. Once n_t reached 10 around year 2000, the consolidation process through mergers increased the HHI from $1,500$ to $2,500$ during the first decade of the 21st century, and then to almost $4,000$ on the back of the 5-to-4 and 4-to-3 mergers.

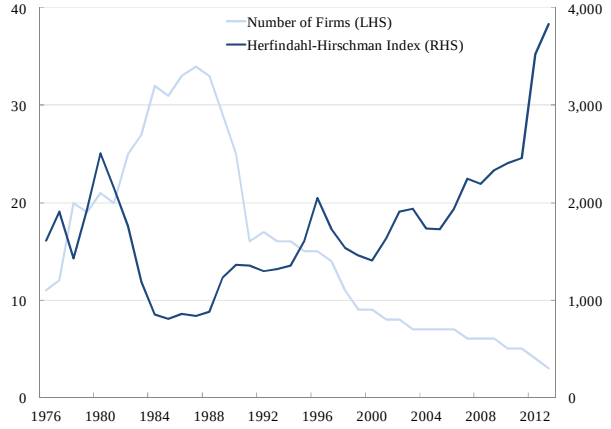
Table 10: Market Shares Before/After Mergers (%)

Year	Target name	Acquiror name	ms^T Before	ms^A Before	$ms^T + ms^A$ Before	$ms^T + ms^A$ After
1982	Memorex	Burroughs	7.83	1.85	9.68	2.73
1983	ISS/Univac	Control Data	0.75	27.08	27.83	19.85
1984	Vertex	Priam	0.93	2.52	3.45	2.78
1988	Plus Dev.	Quantum	0.89	1.41	2.30	4.64
1988	Imprimis	Seagate	13.92	18.16	32.08	29.23
1989	MiniScribe	Maxtor	5.68	4.99	10.68	8.53
1994	DEC	Quantum	1.65	18.60	20.25	20.68
1995	Conner	Seagate	11.94	27.65	39.58	35.41
2001	Quantum	Maxtor	13.87	13.87	27.73	26.84
2002	IBM	Hitachi	13.86	3.64	17.50	17.37
2006	Maxtor	Seagate	8.19	29.49	37.67	35.27
2009	Fujitsu	Toshiba	4.41	10.32	14.72	11.26
2011	Samsung	Seagate	6.89	39.00	45.89	42.82
2012	Hitachi	Western Digital	20.32	24.14	44.46	44.27

Note: ms^T and ms^A denote the target and the acquiring firms' market shares, respectively. For each merger case, "before" refers to the last calendar quarter in which ms^T was recorded separately from ms^A, and "after" is four quarters after "before." Alternative time windows including 1, 8, and 12 quarters lead to similar patterns.

Source: *DISK/TREND* Reports (1977–99), *TRENDFOCUS* Reports (1996–2014), and interviews.

Figure 9: Herfindahl-Hirschman Index (HHI) of the Global HDD Market



Note: The HHI is the sum of the squares of the firm's market shares.

Appendix D Supplementary Materials for Section 4

D.1 Supplementary Materials for Section 4.1

Demand Estimates by Subsample

Our initial demand estimates (Table 3, columns 1, 2, and 3) used the entire sample period, implicitly assuming that the demand function remained constant over time. However, changing uses of digital technology could have altered the consumers' willingness to pay for the same amount of data storage. To investigate this possibility, we estimate our demand model using two subsamples (i.e., the first and the second halves). Table 11 shows that the first-half and the second-half estimates for the main parameter, the price coefficient (α_1), are within the 95% confidence intervals of each other, across all of the three specifications. Thus consumers' valuation for gigabytes of data storage does not exhibit time trend in a statistically significant manner.

Table 11: Demand Estimates by Subsample

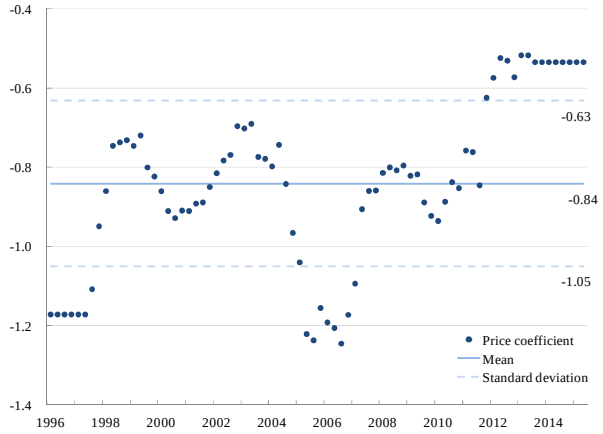
Dependent variable:	(1)		(2)		(3)	
log total EB shipped	OLS		IV-1		IV-2	
Subsample period:	First half	Second half	First half	Second half	First half	Second half
log price per GB (α_1)	$-.8165^{***}$	$-.8594^{***}$	$-.8188^{***}$	$-.8504^{***}$	$-.8959^{***}$	$-.8624^{***}$
	(.0246)	(.0264)	(.0172)	(.0233)	(.0484)	(.0249)
log PC shipment (α_2)	$.8053^{***}$	1.6302^{***}	$.7896^{***}$	1.6191^{***}	.2773	1.6340^{***}
	(.1728)	(.2422)	(.1222)	(.3809)	(.3191)	(.3797)
Constant (α_0)	-1.6405^{***}	-4.3901^{***}	-1.5868^{***}	-4.3337^{***}	.1649	-4.4094^{***}
	(.5863)	(.8718)	(.4102)	(1.3306)	(1.0895)	(1.3275)
Number of observations	39	39	39	39	39	39
Adjusted R^2	.9972	.9746	.9973	.9765	.9966	.9766
First stage regression						
IV for HDD price	–	–	*Disk price*	*Disk price*	*Time trend*	*Time trend*
F-value	–	–	2973.32	536.17	350.51	1056.23
Adjusted R^2	–	–	.9944	.9638	.9346	.9824

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

Figure 10 shows the price-elasticity estimates from a refined version of the estimation by subsample, in which we roll through the full sample with a 12-quarter window. The use of wider windows (i.e., 16 and 20 quarters) produced similar patterns, whereas a narrower window with 8 quarters led to highly volatile estimates.

We see no obvious trend. The exceptions are several low values (i.e., *more* elastic demand) at the beginning and higher values (i.e., *less* elastic demand) at the end, but the first 6 and

Figure 10: Herfindahl-Hirschman Index (HHI) of the Global HDD Market



Note: Each dot represents a 12-quarter rolling estimate of the price coefficient, α_1. This plot visualizes Table 3 (column 4).

the last 6 quarters are carbon copies of the adjacent estimates because one can run less than 78 regressions on a sample of 78 observations. The cyclicality has no obvious explanations either. The IT boom around 2000 coincides with marginally more elastic demand, but no such event accompanied another streak of elastic demand around 2006. Thus we do not see a time trend or systematic fluctuations in our rolling estimates of people's willingness to pay. Nevertheless, we believe it is natural that the demand parameters for high-tech products exhibit some variation across time, and have chosen to use this rolling-estimation version as the baseline for subsequent analysis.

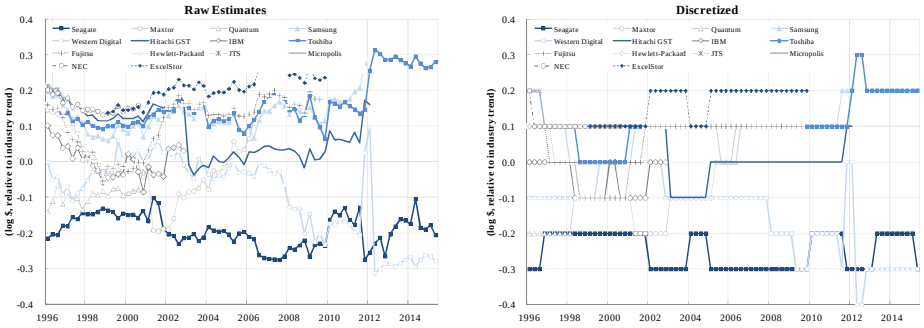
D.2 Supplementary Materials for Section 4.2

Discretization of Productivity Levels

We define the empirical state space by discretizing the levels of firm-specific productivity based on the marginal cost estimates in Section 4.2. Figure 11 (left) plots the trajectories of marginal costs at the firm level between 1996 and 2015. Because the entire industry has experienced a secular trend of cost reduction, we de-trend these estimates and express them relative to the trajectory of Kryder's Law, in the natural logarithm of dollars.

To parameterize the dynamic oligopoly game parsimoniously and keep it computationally tractable, we discretize this relative marginal-cost space as shown in Figure 11 (right). This discretization scheme eliminates small wiggles of productivity evolution but preserves the overall patterns of these firms' relative performances, including their major shifts as well

Figure 11: Marginal Cost Estimates and Their Discretization



Note: The left panel plots our marginal cost estimates. The right panel displays its discretized version.

as leader-follower differences (at least most of the persistent ones). Finer grids resulted in too many zig-zag patterns, frequently amplifying small wiggles which happened to cross the discretization thresholds. More coarse grids tended to eliminate such noises, but the transitions between levels became too infrequent and each of these productivity changes became too impactful in terms of its profit implications via Cournot competition. After experimenting with these alternative grids, we have come to prefer the 0.1 log-dollar grid because it appears to strike the right balance between noise reduction and smooth transitions.

These discretized marginal cost estimates (say, $\overline{mc}_{it}$s) span the state space of firm-specific productivity levels, which will be denoted by $\omega_{it} \in \{\bar{\omega}_1, \bar{\omega}_2, ..., \bar{\omega}_M\}$, where $M = 7$ with our preferred grid. Note that the ranking convention reverses as we redefine marginal costs as productivity levels. That is, a lower marginal cost will be referred to as a high productivity level.

D.3 Supplementary Materials for Section 4.3

Fixed Costs and Accounting Data

We determine the fixed cost of operations and technological catch-up, κ^c, directly from accounting data rather than estimating it along with the three sunk-cost parameters $(\kappa^i, \kappa^m, \kappa^e)$ in Section 4.3, for the following reasons. Our previous experience with the estimation of dynamic games (i.e., Igami 2015, 2016; Igami and Yang 2016) suggests the fixed cost of operations is an order of magnitude smaller than the sunk costs of entry and other major investments (e.g., product and process innovations). Moreover, the fixed-cost estimates tend

to be statistically indistinguishable from zero when sparse data are used, and play a relatively minor role in the overall performance of the dynamic models. Thus, rather than adding κ^c as another parameter to the main estimation procedure, we prefer pinning it down separately from auxiliary data, such as the firms' financial statements.

Accounting data are not always conceptually equivalent to the objects in economic models, as our discussion of profits in Section 4.2 clarifies. But they are nevertheless useful for some purposes, such as fixing the values of a relatively unimportant parameter that cannot be precisely estimated anyway. Our notion of κ^c is something stable over time, and the accounting data on SGA and R&D expenses share this property.

Table 12: Summary Statistics of Accounting Data on Fixed Costs

Variable	Unit of measurement	Number of observations	Mean	Standard deviation	Minimum	Maximum
Fixed cost, κ^c	Million \$	35	1,078	686.7	230.9	2,422
Year, t	Fiscal year	35	2,007	5.419	1,996	2,015
Productivity level, ω_{it}	Levels 1–7	35	4.943	0.996	3	6
Indicator$\{i = Seagate\}$	0 or 1	35	0.428	0.502	0	1
Indicator$\{(i,t) \in Special\}$	0 or 1	35	0.114	0.323	0	1

We estimate κ^c from the financial statements of Seagate Technology and Western Digital between 1996 and 2015. We rely on these firms simply because they are the only publicly traded companies for which systematic records exist. Moreover, they specialize in the manufacturing of HDDs, whereas other survivors such as Hitachi and Toshiba are conglomerates and disclose limited information on HDD-specific activities. The two firms clearly represent a highly selective sample but not a terrible source of information when our only purpose is to capture a ballpark trend in operating costs over two decades.

Table 12 shows summary statistics. Sample size is smaller than 40 (i.e., 2 firms times 20 years) because Seagate became privately owned for financial restructuring in 2000 and its financial statements lost consistency after it went public again. Our main variable is *fixed cost*, which is the sum of SGA and R&D expenses. The right-hand side variable include *year*, *productivity level* (based on the discretized version of our marginal-cost estimates), *Seagate dummy* (the omitted category is Western Digital), and a *special-occasion dummy* (to distinguish abnormal periods for Western Digital when its facilities were hit by a natural disaster).

Table 13 shows the results of OLS regressions. Time trend is positive and statistically significant, whereas productivity level (i.e., control for concurrent firm sizes) is positive but

Table 13: Fixed Cost Estimates from Accounting Data

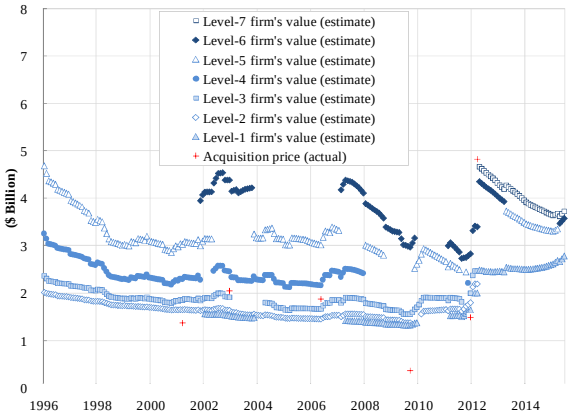
Dependent variable:	(1)	(2)	(3)	(4)
Fixed cost, κ^c	OLS	OLS	OLS	OLS
Year (t)	89.80^{***}	–	61.73^{***}	48.13^{***}
	(11.16)	(–)	(13.95)	(9.81)
Productivity level (ω_{it})	–	540.09^{***}	332.30^{***}	25.96
	(–)	(74.50)	(75.86)	(72.76)
$I\{(i,t) \in Special\}$	–	–	–	728.61^{***}
	(–)	(–)	(–)	(132.69)
$I\{i = Seagate\}$	–	–	–	$1,182^{***}$
	(–)	(–)	(–)	(187.4)
Number of observations	35	35	35	35
Adjusted R^2	0.633	0.603	0.746	.888

imprecisely estimated presumably because of multi-collinearity. Historically, Seagate spent more than Western Digital, but the latter had to spend large sums when it recovered from a flood in Thailand in October 2011. We use predicted fixed costs based on the last (full) specification as $\kappa_t^c(\omega_{it})$ in our main estimation task in Section 4.3.

Implied vs. Actual Acquisition Prices

We conduct a sanity check of fit by comparing predicted enterprise values and the actual acquisition prices. Figure 12 plots our firm-value estimates along the historical path of market structure in the data, and overlays the actual transaction prices in the six merger cases from *Thomson*'s financial data (marked by red crosses). Because target firms' stand-alone values underpin their equilibrium acquisition prices in our model, comparison of the estimated values and the actual acquisition prices provides a ballpark assessment of the fit in terms of dollar values. In at least half of the cases, each of the acquisition prices is located close to the estimated value of firms with the corresponding productivity level (1, 2, 3, or 4) and stays within the range of the focal level and its adjacent level. Thus we regard the estimated model as a reasonable benchmark with which we may compare our counterfactual simulation to assess the impacts of a hypothetical merger policy.

Figure 12: Firm Value Estimates and Actual Acquisition Prices



Note: Red crosses represent the actual acquisition prices in the six merger cases from *Thomson* database. The other seven markers represent our estimates of equilibrium firm values along the historical path of market structure in the data.

References

[1] Aghion, Philippe, Nick Bloom, Richard Blundell, Rachel Griffith and Peter Howitt (2005) "Competition and innovation: An inverted-U relationship," *Quarterly Journal of Economics*, 120 (2): 701–728.

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

[3] Arcidiacono, Peter, Patrick Bayer, Jason R. Blevins, and Paul B. Ellickson (2015) "Estimation of Dynamic Discrete Choice Models in Continuous Time with an Application to Retail Competition," *Review of Economic Studies*, forthcoming.

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

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

[6] Baron, David P., and John A. Ferejohn (1989) "Bargaining in Legislatures," *American Political Science Review*, 83 (4): 1181–1206.

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

[8] Berry, Steven T., and Ariel Pakes (1993) "Some Applications and Limitations of Recent Advances in Empirical Industrial Organization: Merger Analysis," *American Economic Review Papers and Proceedings*, 83(2): 247–252.

[9] Collard-Wexler, Allan (2013) "Demand Fluctuations in the Ready-Mix Concrete Industry," *Econometrica*, 81 (3): 1003–1037.

[10] ———, Gautam Gowrisankaran, and Robin S. Lee. 2014. "'Nash-in-Nash' Bargaining: A Microfoundation for Applied Work," NBER Working Paper, No. w20641.

[11] Crawford, Gregory S., and Ali Yurukoglu (2012) "The Welfare Effects of Bundling in Multichannel Television Markets," *American Economic Review*, Vol. 102, pp. 643-85.

[12] Dasgupta, Partha and Joseph Stiglitz (1980) "Industrial Structure and the Nature of Innovative Activity," *Economic Journal*, Vol. 90, No. 358, pp. 266-293.

[13] Demsetz, Harold (1973) "Industry structure, market rivalry, and public policy," *Journal of Law and Economics*, April, 1–90.

[14] Deneckere, Raymond and Carl Davidson (1985) "Incentives to Form Coalitions with Bertrand Competition," *RAND Journal of Economics*, Vol. 16, pp. 473-86.

[15] Egesdal, Michael, Zhenyu Lai, and Che-Lin Su (2015) "Estimating dynamic discrete-choice games of incomplete information," *Quantitative Economics*, 6: 567–597.

[16] Entezarkheir, Mahdiyeh, and Saeed Moshiriz (2015) "Merger Induced Changes of Innovation: Evidence from a Panel of U.S. Firms," mimeo.

[17] Ericson, Richard, and Ariel Pakes (1995) "Markov-Perfect Industry Dynamics: A Framework for Empirical Work," *Review of Economic Studies*, 62 (1): 53–82.

[18] Farrell, Joseph and Carl Shapiro (1990) "Horizontal Mergers: An Equilibrium Analysis," *American Economic Review*, Vol. 80, pp. 107-26.

[19] Federal Trade Commission (2013) "Horizontal Merger Investigation Data: Fiscal Years 1996–2011."

[20] Gilbert, Richard J., and Hillary Greene (2015) "Merging Innovation into Antitrust Agency Enforcement of the Clayton Act," *George Washington Law Review*, 83: 1919–1947.

[21] ——— and David Newbery (1982) "Preemptive Patenting and the Persistence of Monopoly," *American Economic Review*, 72(2): 514–26.

[22] Goettler, Ronald, and Brett Gordon (2011) "Does AMD spur Intel to innovate more?" *Journal of Political Economy*, 119(6): 1141–200.

[23] Gowrisankaran, Gautam (1995) "A Dynamic Analysis of Mergers," Ph.D. dissertation, Yale University.

[24] ——— (1999) "A Dynamic Model of Endogenous Horizontal Mergers," *RAND Journal of Economics*, Vol. 30, pp. 56-83.

[25] Ho, Kate (2009) "Insurer-Provider Networks in the Medical Care Market," *American Economic Review*, Vol. 99, pp. 393-430.

[26] Igami, Mitsuru (2015) "Estimating the Innovator's Dilemma: Structural Analysis of Creative Destruction in the Hard Disk Drive Industry, 1981–1998," *Journal of Political Economy*, forthcoming.

[27] ——— (2016) "Industry Dynamics of Offshoring: The Case of Hard Disk Drives," mimeo.

[28] ——— and Nathan Yang (2016) "Unobserved Heterogeneity in Dynamic Games: Cannibalization and Preemptive Entry of Hamburger Chains in Canada," *Quantitative Economics*, 7 (2): 483–521.

[29] Iskhakov, Fedor, John Rust, and Bertel Schjerning (2014) "The Dynamics of Bertrand Price Competition with Cost-Reducing Investments," manuscript, Georgetown University.

[30] ——— (2016) "Recursive Lexicographical Search: Finding All Markov Perfect Equilibria of Finite State Directional Dynamic Games." *Review of Economic Studies*, 83 (2): 658-703.

[31] Iskhakov, Fedor, Jinhyuk Lee, John Rust, Bertel Schjerning, and Kyoungwon Seo (2016) "Comment on 'Constrained Optimization Approaches to Estimation of Structural Models'," *Econometrica*, 84 (1): 365–370.

[32] Jeziorski, Przemyslaw (2014) "Empirical Model of Dynamic Merger Enforcement: Choosing Ownership Caps in U.S. Radio," mimeo.

[33] Kim, Myongjin (2015) "Strategic Responses to Used Goods Markets: Airbus and Boeing." Manuscript, University of Oklahoma.

[34] Kreps, David M. and Jose A. Scheinkman (1983) "Quantity precommitment and Bertrand competition yield Cournot outcomes," *The Bell Journal of Economics*, Vol. 14, No. 2, pp. 326–37.

[35] Lamoreaux, Naomi R.. 1985. *The Great Merger Movement in American Business, 1895-1904* (New York: Cambridge University Press). Reissued in paperback, 1988.

[36] Marshall, Guillermo, and Álvaro Parra (2015) "Mergers in Innovative Industries," mimeo.

[37] McFarland, Henry, and Philip Nelson (2008) "Failing Firms and Declining Industries," *Issues in Competition Law and Policy*, 3: 1691–1716.

[38] Mermelstein, Ben, Volker Nocke, Mark A. Satterthwaite, and Michael D. Whinston (2014) "Internal Versus External Growth in Industries with Scale Economies: A Computational Model of Optimal Oligopoly Policy," mimeo.

[39] Miller, Nathan H., and Matthew C. Weinberg (2015) "Mergers Facilitate Tacit Collusion: Empirical Evidence from the U.S. Brewing Industry," Manuscript, Georgetown University.

[40] Nevo, Aviv (2000) "Mergers with Differentiated Products: The Case of the Ready-to-Eat Cereal Industry," *RAND Journal of Economics*, Vol. 31, No. 3, pp. 395-421.

[41] Okada, Akira (1996) "A Noncooperative Coalitional Bargaining Game with Random Proposers," *Games and Economic Behavior*, 16: 97–108.

[42] Ozcan, Yasin (2015) "Innovation and Acquisition: Two-Sided Matching in M&A Markets," mimeo.

[43] Pakes, Ariel (1986) "Patents as options: Some estimates of the value of holding European patent stocks," *Econometrica*, 54 (4): 755–784.

[44] ———, Michael Ostrovsky, and Steven Berry (2007) "Simple estimators for the parameters of discrete dynamic games (with entry/exit examples)," *RAND Journal of Economics*, 38 (2): 373–399.

[45] Perry, Martin K. and Robert H. Porter (1985) "Oligopoly and the Incentive for Horizontal Merger," *American Economic Review*, Vol. 75, pp. 219-27.

[46] Pesendorfer, Martin, and Philipp Schmidt-Dengler (2008) "Asymptotic Least Squares Estimators for Dynamic Games," *Review of Economic Studies*, 75 (3): 901–928.

[47] Qiu, Larry and Wen Zhou (2007) "Merger Waves: a Model of Endogenous Mergers," *RAND Journal of Economics*, Vol. 38, pp. 214-226.

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

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

[50] Salant, Stephen W., Sheldon Switzer, and Robert J. Reynolds (1983) "Losses from Horizontal Merger: The Effects of an Exogenous Change in Industry Structure on Cournot-Nash Equilibrium," *Quarterly Journal of Economics*, Vol. 98, pp. 185-199.

[51] Scherer, Frederic M. (1965) "Firm Size, Market Structure, Opportunity, and the Output of Patented Inventions." *American Economic Review*, 55(5): 1097–1125.

[52] Schumpeter, Joseph A. (1942) *Capitalism, Socialism, and Democracy*, Harper & Brothers, NY.

[53] Scotchmer, Suzanne (2004) *Innovation and Incentives*, MIT Press, Cambridge, MA.

[54] Seim, Katja (2006) "An empirical model of firm entry with endogenous product-type choices," *RAND Journal of Economics*, 37 (3): 619–640.

[55] Selten, Richard (1973) "A Simple Model of Imperfect Competition, where 4 Are Few and 6 Are Many," *International Journal of Game Theory*, 2(1): 141–201.

[56] Stahl, C. Jessica (2011) "A Dynamic Analysis of Consolidation in the Broadcast Television Industry," mimeo.

[57] Stigler, George J. (1950) "Monopoly and Oligopoly by Merger," *American Economic Review*, Papers and Proceedings, 40 (2): 23–34.

[58] Sutton, John (1998) *Technology and Market Structure: Theory and History*. Cambridge, MA: The MIT Press.

[59] Takahashi, Yuya (2015) "Estimating a War of Attrition: The Case of the US Movie Theater Industry," *American Economic Review*, 105 (7): 2204–41.

[60] Tirole, Jean (1988) *The Theory of Industrial Organization*, MIT Press, Cambridge, MA.

[61] Werden, Gregory J. and Luke M. Froeb (1994) "The Effects of Mergers in Differentiated Products Industries: Logit Demand and Merger Policy," *Journal of Law, Economics, & Organization*, Vol. 10, No. 2, pp. 407-426.

[62] Williamson, Oliver E. (1968) "Economies as an Antitrust Defense: The Welfare Tradeoffs," *American Economic Review*, Vol. 58, No. 1, pp. 18-36.