Extending the Frontier: A Structural Model of Investment and Technological Competition in the Supercomputer Industry

João Macieira*

November 6, 2006

Abstract

How does technological progress depend on competition through investments in innovation? This paper proposes and estimates a structural model of dynamic investment and competition in the supercomputer industry to evaluate the dependence of technological innovation on market structure. Building on recent advances in the structural estimation of investment games, I develop a tractable yet dynamic model of the benefits and costs of innovation, including the dependence of innovation cost on the current technological position of the firm and accounting for the dynamic benefits of technological leadership. This model identifies the innovation incentives for all firms in the industry, and yields an estimable model of the structure of profits and investment costs. Taking advantage of a novel and unique panel data set comprising nearly all supercomputer purchases over a 16-year period, the approach builds upon and extends recent advances in the estimation of dynamic investment in the presence of strategic interactions. The estimates facilitate counterfactual comparisons of how the evolution of the maximal computing speed supplied in the supercomputer industry differs under different market structures. Consistent with the importance of a technology "selection effect" (Aghion et al, 2001), increased levels of competition are associated with a higher rate of innovation in the supercomputer industry. Increased competition is also associated with increased welfare, but the marginal increase in welfare is decreasing in the number of competitors.

*Department of Economics, Northwestern University. I am grateful to my committee Aviv Nevo, Robert Porter, Igal Hendel and Scott Stern for their guidance and support throuhgout this research. I would also like to thank Shane Greenstein, Andrew Sweeting, Elie Tamer, Daniel Ackerberg, Jack Dongarra, Federico Rossi, Sara Ferreira Filipe, Luis Vasconcelos, Adeline Delavande and Carlos Madeira for very helpful discussions and comments. Special thanks are also due to John McCallum for extremely helpful references on suplementary data sources to the TOP500 dataset. Financial support from the Fundação Calouste Gulbenkian and Fundação para a Ciência e Tecnologia (FCT) are also greatfully acknowledged. Comments welcome at j-macieira@northwestern.edu. All errors are my own.

1 Introduction

Technological innovation plays a central role in economic growth. However, there is no consensus on whether competition helps technological innovation. In addition, it is hard to quantify the gains (or losses) from increased competition in high-technology industries. In this paper, I aim to fill this gap. I focus on one high-technology industry - supercomputers - and evaluate the evolution of the best computational speed available in the supercomputer market under three different market structures: monopoly, duopoly and a three-firm market. In addition, I quantify the temporal evolution of consumer welfare and firm profits under these three different industry sizes.

I propose a structural model where firms invest in innovation and set prices. Using data from the supercomputer industry, I estimate the model parameters with a variant of the two-step method of Bajari, Benkard and Levin (2006). I then use the estimated parameters to simulate the outcomes of three different market structures. This is done in two steps: (i) computation of the Markov-Perfect Nash Equilibrium (MPNE) of the game under monopoly[1], duopoly and three-firm market.; and (ii) comparison of consumer welfare, profits, and the maximal available computational speed in the supercomputer industry implied by these industry structures for a period of five consecutive years.

I find that increased levels of competition have positive effects on the evolution of supercomputer technology. In line with Aghion, Harris, Howitt and Vickers (2001), increases in market competition induce more innovation due to the "selection effect". That is, the incremental payoff from innovating is higher when a firm is in "neck-and-neck" competition with technologically similar rivals. Firms therefore innovate primarily to escape competition with "neck-and-neck" rivals. I also find that increased competition increases welfare in the supercomputer market, even though the marginal increase in welfare is decreasing in the number of competitors. Finally, I find that fixed innovation costs are the most important component of technology investment expenditures. Both the firm and its rivals' technological states impact fixed costs considerably. In particular, this component of the innovation costs is decreasing in the aggregate frontier of a firm's most advanced rivals. Even though fixed innovation costs are decreasing on the firm's technological state, the technological leader has a remarkable extra fixed cost on expanding its frontier.

My model builds on the refinements proposed by Doraszelski and Satterthwaite (2005) to the theoretical model of Ericson and Pakes (1995). One of the major challenges is how to deal with multi-product firms. I build on the model proposed by Nevo and Rossi (2006), in which a markup-adjusted inclusive value is included as state variable. This allows me to control for product market profits while keeping a computationally manageable state space dimension. A second difficulty is how to deal with nonstationary states and controls. I follow the methods of variable scaling in the growth literature surveyed in Stockey and Lucas (1989) to guarantee stationary measures of firm technological state and innovation investment.

[1] In the monopoly case, the MPNE collapses into a single-firm investment optimization problem.

To estimate the model I use a variant of the two step method proposed by Bajari, Benkard and Levin (2006). In the first step, I recover the profit function parameters from demand and supply estimation. In addition, I estimate policies and state transitions nonparametrically as functions of observable states. In the second step, I forward simulate the continuation values of firms using the policy and transition estimates from the first stage. Instead of using the continuation values in the inequality sampling method proposed by Bajari et al. (2006), I include them in an estimation procedure where I use firm optimality conditions directly. I am able to separate the benefits from innovation from its costs by using the variations on the number of a firm's technologically similar rivals seen in the data. A higher number of similar rivals does not impact the costs of innovation but increases the incremental payoff from innovating due to the "selection effect". Therefore, the positive correlation between investment rates and number of technologically similar rivals seen in the data successfully identifies the benefits from innovation.

The understanding of the forces driving innovation and the measurement of competition effects require focus on a specific industry. Ideally, one would want to analyze an industry whose technological advances influence the evolution of other industries' technologies. Bresnahan and Trajtenberg (1995) document 'too late, too little' innovation in industries producing this type of technology - named General Purpose Technology (GPT). Moreover, the positive feedbacks of advances in GPTs are usually dispersed through the economy, and can therefore work as a trigger for economic growth. One of these industries, supercomputers, is particularly well suited to study the technological evolution patterns under different market structures for several reasons. First, there is a standard measure for supercomputer quality - the LINPACK Benchmark. A supercomputer, commonly mentioned as High Performance Computer (HPC), is a general purpose computer that is faster than commercial competitors and has sufficient central memory to compute problem sets of general scientific interest. The technological state of a firm in this industry can be measured by the maximal supercomputer quality that the firm has ever produced. Second, the technological evolution patterns of the firms in this industry exhibit evidence of "racing behavior", where players aim to leapfrog opponents as documented by Khanna (1995). Therefore, the supercomputer industry provides an almost ideal ground to study innovation behavior under oligopoly. Third, supercomputer technology brings several important benefits that are only partially considered by manufacturers, thus justifying governmental intervention. Supercomputers are crucial to advances in many sciences and play an important role on production of other goods (e.g. cars, airplanes, medicines, intelligence creation and national defense).

A recent body of the industrial organization literature studies the incentives for technical progress in oligopolistic industries. (e.g., Athey and Schmutzler (2001), Aghion et al. (2001), Scotchmer(1991,2004)). More recently, Segal and Whinston (2004) analyze the effects of specific antitrust policies in industries where innovation is of central importance to competition. Most of these studies restrict attention to duopoly games, or abstract from important dynamic strategic considerations of players (e.g. dependence of innovation decisions on the firms' technological states). I introduce a model where such restrictions are relaxed.

This paper also extends the endogenous quality choice literature, which has rarely focused on dynamic competition models[2]. Carranza (2006) proposes a dynamic model of product innovation in the digital camera market where the quality of new products is optimally decided by each firm. However, the impact of firms' quality decisions in the state space is assumed to be negligible in Carranza's model. This assumption cannot be considered for a model of competition through innovation, since it defeats the purpose of relating the individual firms' actions to observed technological states. In contrast, my modeling of firms' optimal investment on technological innovation explicitly accounts for both firm strategies' dependence on the state variables and its impact on the state space.

The rest of the paper is organized as follows. In section 2 I describe the supercomputer industry, mentioning its key differences compared to other high-end computers. In Section 3 I present the dynamic oligopoly model, describing firm behavior on pricing and innovation investment. The details about the data and the empirical strategy for model estimation are presented in section 4. Section 5 addresses the estimation results, while section 6 describes both the policy experiments and the simulation results. Section 7 concludes with results and contributions summary, as well as possible extensions for the paper. All proofs, derivations and computational details can be found in the Appendixes.

2 The supercomputer industry

A supercomputer refers to "those computing systems (hardware, systems software, and applications software) that provide close to the best currently achievable sustained performance on demanding computational problems"[3]. Supercomputers are very expensive durables whose expected useful lifetime is five years. Also known as high-performance computer (HPC), its key components are processors, memory and interconnect bandwidth. The processors are assigned the tasks of performing the instructions programed in the supercomputer (e.g., simulation of a nuclear explosion). In order to temporarily store data or results for processing in intermediate program instructions, all supercomputers must contain a memory sector which must be connected with the processors. An interconnect bandwidth is required to control the traffic of information between the memory sector and the processors. The observable quality of a supercomputer can be measured by the total floating-point operations per second (FLOPS). Even though there are several benchmarks for speed measurement of a supercomputer, the consensus quality measure is the Linpack Benchmark. This is the maximum possible computation speed that can be achieved by the supercomputer when it is instructed

[2]Some of the most recent research on static endogenous quality choice can be found in Mazzeo(2002) and Seim (2002).
[3]Source: "Getting Up to Speed: the Future of Supercomputing" (2005).

to solve a system of linear equations with 1000 equations and 1000 unknowns[4].

Despite their similarities, there are two main differences between a supercomputer and a mainframe. The first is that, upon delivery and installation, the supercomputer is configured for selected purposes, chanelling all its power to execute these tasks as fast as possible. High performance computers are assigned very specific tasks (e.g., simulation of car acidents for new vehicle models, simulation of nuclear explosions) to be performed one at a time, whereas a mainframe uses its power to performs several programs simultaneously. The second main difference between a supercomputer and a mainframe is the system architecture. High performance computers are designed to execute specific programs as efficiently as possible, implying a deviation from the regular mainframe architecture designed for simultaneous task processing.

There are two main types of supercomputer architectures: vector systems (VS) and massive paralell processing systems (MPPS). The vector systems are characterized by special central processors (which is typically produced by the manufacturing firm), along with bult-in memory chips and interconnect bandwidth. Processing takes place by seeding in several tasks at the same time. The processor is specially designed to interpret the vector of tasks and to perform instructions simultaneously. This technique, initially proposed to the late supercomputer engineer Seymour Cray, is know as vector processing. The VS are known for its computing efficiency (i.e., high computation speed due to optimized combination of processors, memory and bandwidth), as well as for acceptable instalation area requirements, reasonable power consumption and maintenace costs. However, it is usually expensive, typically tens of millions of dollars. This is because of the high costs required to build special purpose processors.

This description sharply constrasts with the one for MPPS. Instead of special purpose processor, manufacturers of MPPS use off-the-shelf processors to build a processing cluster. A similar procedure is followed for building the memory and bandwidth sectors of the supercomputer. As such, the production of MPPS is not demanding in R&D efforts. For this reason, this type of system is usually not as expensive as VS. However, MPPS have its own caveats. First, it is not as efficient as VS. It relies on simple clustering of microprocessors along with memory and bandwidth components, while in VS these components are conceived in order to maximize computing efficiency. Second, despite its relatively cheap price, MPPS require far more instalation space, maintenace costs and power consumption.

Despite these disadvantages, MPPS have become popular in supercomputing since the early 1990's for two reasons. First, the prices of off-the-shelf microprocessors have fallen dramatically over time, while its processing power has steadily increased over time. Second, despite the mild progress on bandwidth and memory technology, there was enough scope for increases in MFLOPS by adding extra microprocessors, for a given combination of memory and bandwidth. This advantage over VS is no longer available, as the

[4] A comprehensive discussion on the Linpack Benchmark can be found in Dongarra (2006), and Dongarra, Luzczek and Petitet (2001)

current pace of memory and bandwidth advances can no longer sustain gains in speed from increasing the processor power. The most up-to-date supercomputer models aim to combine the advantages of both types of architectures.

The supercomputer industry is characterized by few units being sold and considerable revenues. The joint revenue of the top nine competitors in this industry was $4655 million in 1997, reaching its peak in 2000 with $6083 million[5]. The vast majority of buyers are governmental institutions ,which account for at least 70% of total supercomputer sales since 1990. Since the rise of this industry in 1953, several firms have entered and exited the market. Most of the entry and exit is due to short-lived competitors, while most of the important players in this market remained active. The most relevant competitors who kept their importance over time have been Cray Inc., Hewlett-Packard, IBM, NEC, Fujitsu, Hitachi, SGI, and Sun Microsystems. The benefits of being among the best HPC suppliers are not restricted to product market profits. Innovations in supercomputer technology are often used for improvements in other high-tech products the firm may be producing. For example, advances in computational speed for HPC have been incorporated in personal computers during the last ten years. Nonsuprisingly, except for Cray, all other supercomputer sellers have other product lines.

Table I provides yearly statistics on the maximum quality ever produced (henceforth *firm technological frontier*) for the most important firms in the market. Several important aspects can be inferred. First, significant increases on the maximal computing speed in the industry tend to be followed by periods of no expansion. In general, the more significant this expansion is, the more years are necessary to beat that record. Second, the data indicates that once a firm reaches quality leadership, it will not expand its frontier significantly (if at all) unless a rival leapfrogs upon it. Therefore, the incentives on firms to invest on frontier expansion are contingent on the firms' frontier position. Finally, firms tend to expand their frontiers more significantly over time. This suggests that firms can benefit from the technological advances of their rivals. Therefore, since the technological advances of a firm leverage rival innovations, firms are likely to underinvest.

Unlike frequently purchased items, supercomputers are acquired in small quantities by procurement. Potential buyers may either choose among existing models (henceforth *"off-the-shelf" HPCs*) or order a machine whose computing power is suited for her specific needs (henceforth *custom supercomputers*). There are some differences between buying an existing model and ordering a custom supercomputer. If the consumer is interested on ordering a custom model, potential manufacturers are called to submit proposals, indicating the prototype details and the price to be charged. Upon testing the prototypes, the buyer chooses the supplier and makes a public announcement of the winning proposal details. For the case of off-the-shelf models, the potential buyer contacts the supercomputer suppliers, in an attempt to obtain discounts over the list prices. After checking if the available models suit his computing needs, the consumer decides whether to place an

[5]Source: IDC 2002

order. Discounts over list price are frequent, but depend on the client characteristics. According to the available information on discounting practices in this industry, discounts for private and non-government clients are in the range of 10%-20% over the list price. Government-related institutions are tipically given a discount rate between 20%-30%[6]. Like other technological goods, supercomputers have experienced sharp decreases in both nominal and real price/quality ratio.

TABLE I

TECHNOLOGICAL FRONTIER STATISTICS FOR THE MOST IMPORTANT FIRMS (IN RMAX GIGAFLOPS)

Year	Cray	HP	IBM	NEC	Fujitsu	Hitachi	SGI	Intel	Sun	Industry Max.
1990	2.171	-	0.54	23.2	4	1.817	-	2.6	-	23.2
1991	2.171	-	0.54	23.2	4	1.817	-	13.9	-	23.2
1992	13.7	-	1.457	23.2	4	1.817	-	15.2	-	23.2
1993	13.7	1.6	5.8	23.2	124	7.4	1.284	143.4	-	143.4
1994	100.5	3.306	66.3	23.2	124	27.5	4.142	143.4	-	143.4
1995	100.5	7.408	88.4	60.72	124	28.4	26.653	143.4	17.91	143.4
1996	341	15.01	88.4	66.53	229	368.2	341	143.4	17.91	368.2
1997	815	51.3	151.8	244	229	368.2	815	1338	26.45	1338
1998	891	51.3	547	244	229	368.2	891	1338	272.1	1338
1999	1166	189.3	2144	244	492	873	1166	2379	420.44	2379
2000	1166	189.3	4938	303	886	1035	1166	-	420.44	4938
2001	1166	431.7	4938	1192	886	1709.1	1166	-	420.44	4938
2002	1166	2916	7304	35860	5406	1709.1	1166	-	1226.4	35860
2003	2932.9	8633	7304	35860	5406	1709.1	11652	-	1226.4	35860
2004	5895	8633	70720	35860	8728	3319	51870	-	1439	70720
2005	36190	8633	280600	35860	8728	3319	51870	-	3146	280600

The scope of application of supercomputers is immense. According to the classification available at the supercomputing rating organization TOP500, there are 26 application areas for supercomputers. Some representative examples application areas are Aerospace, Automotive, Finance, Defense, Geophysics, Semiconductors, Weather and Climate Research, and Telecommunications.

[6] As an example, the latest available contract between Cray and the General Services Administration (GSA) specified a discount rate of about 22% for several supercomputer models and related components.

3 The model

In this section I present a model of firm behavior in the supercomputer industry. Time is assumed discrete with an infinite horizon, and indexed by $t \in 1, 2, ..., \infty$. At each period, multiproduct firms simultaneously decide on innovation investment and retail prices. Investment is defined as being the number of computing speed units (measured in Gigaflops according to the Linpack Benchmark) added to the maximal computing speed that the firm has ever produced (henceforth *firm frontier*). In what follows, I build on the work of Doraszelski and Satterthwaite (2005), who refine the theoretical framework of industry dynamics pioneered by Ericson and Pakes (1995).

Observable states. I assume that all payoff-relevant features of firms can be enconded into a state vector. All firms observe a demand state, M_t, and a pair of firm-specific states, κ_{ft} and ν_{ft}, for all $f = 1, ..., F$ players. κ_{ft} (henceforth *firm rank*) represents the ratio of firm f's frontier over the best frontier available in the industry at time t. Ideally, I would like to consider the firm frontier itself as state variable. Unfortunately, the fact that this variable grows without bound makes its choice as state variable impractical. Instead, I consider a scaled version of this variable, κ_{ft}, by following the variable rescaling methods surveyed by Stockey and Lucas(1989)[7]. Letting h_{ft} be the maximal computing speed that the firm has ever produced up to time t, κ_{ft} is defined as

$$\kappa_{ft} \equiv \frac{h_{ft}}{\max\limits_{i=1,\ldots,F} \{h_{it}\}}$$

I denote $\mathbf{s}_t \in S$ as the vector of all observable states at time t.

Firms. Let $\mathbf{p}_{ft}$ denote the vector of prices charged by firm f at time t for its off-the-shelf supercomputers[8]. In addition, I denote I_{ft} as the investment rate of firm f, defined as the ratio between investment and the firm frontier at time t[9]. At the time of the investment and pricing decisions, each firm observes a pair of private information shocks: a shock in fixed innovation costs, ε_{1ft},, and a shock in marginal innovation costs ε_{2ft}. Therefore, firm f's state space is $S \times \Gamma_{1f} \times \Gamma_{2f}$, where Γ_{if} corresponds to the space of realizations of ε_{ift}.

[7]As explained later in the paper, this scaled variable embodies the dependence of firm frontier on innovation while keeping stationarity.

[8]Each cell of the vector corresponds to the price charged for a given supercomputer model at time t. I am not considering custom supercomputers introduced at time t for two reasons. First, custom models are not part of the firm's product portfolio except after its introduction at time t. These are only considered at period $t+1$ if they are available for sale to other customers. Second, I am able to control for the benefits outside off-the-shelf profits with a function described below.

[9]An alternative would be to consider investment as defined early in this section. However, the data on firm maximal computing speed increases is nonstationary: firms tend to expand their frontiers more significantly over time. Considering investment rate rather than investment successfully circumvents this problem.

Before defining each firm's intertemporal optimization problem, I impose the following assumptions:

Assumption A1 (Markovian pure strategies): In equilibrium, all players' choices are deterministic functions of payoff-relevant information.

Formally, this corresponds to a map $\tilde{\sigma}_f : S \times \Gamma_{1f} \times \Gamma_{2f} \longrightarrow (\mathbf{p}_f, I_f)$, for any firm f. In what follows, I assume $\Gamma_{if} = \mathbb{R}, \forall f = 1, ..., F$ and $\forall i = 1, 2$. That is, shocks in innovation costs can be negative. This is, since innovation may be be less costly due to outside factors (e.g. governmental subisidies for certain innovation projects). A firm's intertemporal optimization problem can be written in recursive form, for any profile $\tilde{\sigma} = (\tilde{\sigma}_1, ..., \tilde{\sigma}_F)$ of Markovian strategies.

Assumption A2: Private shocks on both fixed and marginal innovation costs are independently and identically distributed Normal over time and players. The joint distribution of these variables, $F_{\varepsilon_1, \varepsilon_2}$, is common knowledge to all players.

This assumption is motivated by computational concerns. Allowing for serial correlation implies a significant (and unaffordable) increase in computational burden. Of course, alternative distributions could instead be considered. The normality assumption is imposed for simplicity and computational convenience.

Assumption A3: For every firm f, the private information states ε_{1ft} and ε_{2ft} are assumed independent of observed states $\mathbf{s}_t$.

I will use this assumption below not only to reduce the state space, but also to provide a clear equilibrium concept for the game. Assumptions A1-A3 allow me to define firm's problem in a tractable way. The flow payoff of firm f at time t is defined as

$$\pi_f(\mathbf{p}_t, I_{ft}, \mathbf{s}_t, \varepsilon_{ft}) \equiv \Pi_f(\mathbf{p}_t, \mathbf{s}_t) + \Upsilon_f(\mathbf{s}_t) - C(I_{ft}, \mathbf{s}_t, \varepsilon_{ft})$$

where $\varepsilon_{ft} = (\varepsilon_{1ft}, \varepsilon_{2ft})$, $\mathbf{s}_t \equiv (M_t, \hat{s}_t)$, $\hat{s}_t \equiv (\nu_i, \kappa_i)_{i=1}^{N}$, $\mathbf{p}_t \equiv \{\mathbf{p}_{ft}\}_{f=1}^{F_t}$ is a tuple of prices charged by the F firms, $\Pi_f(\mathbf{p}_t, \mathbf{s}_t)$ corresponds to firm f's flow market profits, $\Upsilon_f(\mathbf{s}_t)$ is a function representing the benefits (other than flow profits) for firm f from the current state of the industry $\mathbf{s}_t$. As described in more detail below, these benefits outside profits may include spillovers of frontier expansions to other lines of product within the firm. $C(I_{ft}, \mathbf{s}_t)$ represents innovation costs. The firm is assumed to decide investment and retail prices in order to maximize the expected sum of discounted payoffs. Neither the researcher nor individual firms observe the private shocks of other firms. Therefore, the Bellman equation of the firm is

$$V_f(\mathbf{s},\varepsilon_f)=\int_{\varepsilon_{-f}} \underset{\tilde{\sigma}_f(\mathbf{s},\varepsilon_f)}{Max}\left\{\pi_f(\sigma_f(\mathbf{s},\varepsilon_f),\mathbf{s},\varepsilon_f)+\beta E_{\mathbf{s},,\sigma_f,\sigma_{-f}}\left[V_f(\mathbf{s}',\varepsilon_f')|\mathbf{s},\tilde{\sigma}_f(\mathbf{s},\varepsilon_f),\tilde{\sigma}_{-f}(\mathbf{s},\varepsilon_{-f})\right]\right\}dF(\varepsilon_{-f})$$

where $\tilde{\sigma}_f(\mathbf{s},\varepsilon_f) \equiv (p_f(\mathbf{s}_t,\varepsilon_f), I_f(\mathbf{s},\varepsilon_f))$ represents a Markovian strategy for firm f, and $E_{\mathbf{s},\sigma_f,\sigma_{-f}}[.]$ denotes firm f's expectations conditional on all firms choosing Markovian strategies, and on observable states. However, ε_f is not observed by the researcher. Under assumptions A2 and A3, I can restrict attention to a simplified version of the firm's problem where all private information is integrated out[10]. Therefore, the firm's expected discounted payoffs conditional only on observable states are given by

$$V_f(\mathbf{s})=\int_{\varepsilon} \underset{\tilde{\sigma}_f(\mathbf{s},\varepsilon_f)}{Max}\left\{\pi_f(\sigma_f(\mathbf{s},\varepsilon_f),\mathbf{s},\varepsilon_f)+\beta\int V_f(\mathbf{s}')dP(\mathbf{s}'|\mathbf{s},\tilde{\sigma}_f(\mathbf{s},\varepsilon_f),\tilde{\sigma}_{-f}(\mathbf{s},\varepsilon_{-f}))\right\}dF(\varepsilon)$$

where $P(\mathbf{s}'|\mathbf{s},\tilde{\sigma}_f(\mathbf{s},\varepsilon_f),\tilde{\sigma}_{-f}(\mathbf{s},\varepsilon_{-f}))$ corresponds to the transition of states conditional on current states and actions of all players. The next subsections address the details on product market profits, benefits from industry states (excluding profits), investment costs, state transitions and the equilibrium concept considered for this dynamic game.

3.1 Product Market Profits

In this subsection I model product market profits. In what follows, I impose the following assumption

Assumption A4: Prices do not influence the evolution of both observable and privately known states.

The modeling of product market profits proceeds in three steps. First, I model the demand for supercomputers. Second, I assume price competition in a differentiated products market to model supply and establish its equilibrium properties. Finally, I derive the average quality state ν_{ft} and write each firm's market profit as a function of observed states.

3.1.1 Demand

I follow the characteristics-based discrete choice approach described in McFadden (1981). Let J_t be the number of supercomputer models available at time t. The utility function for a buyer r interested in acquiring supercomputer j at time t is assumed to be

[10] Note that Assumptions A2 and A3 are not necessary for this result. However, the computation of the expected discounted payoffs is considerably simpler under these assumptions.

$$U_{rjt} = \gamma x_j + \delta\kappa_{f(j)t} + \lambda MPP_j - \alpha p_{jt} + \xi_{jt} + \epsilon_{rjt}$$

where x_j is the observed quality of product j, defined as the computing speed measured in Teraflops (1000 Gigaflops) according to the LINPACK Benchmark, MPP_j is a dummy variable which equals unity if supercomputer j belongs to the family of massive paralell processing systems, p_{jt} is the price of supercomputer j at time t in \$1M, ξ_{jt} regards its unobserved quality, and $\kappa_{f(j)t}$ corresponds to the computing speed ranking of the firm which commercializes product j at time t[11]. The richness of the top500 data allows for additional characteristics to be included in the mean utility (e.g., number of processors, operative system and interconnect type). There are, however, two reasons why a more parsimonious model is preferred. First, most of these variables are strongly correlated with the observed quality measure - Rmax -, which raise colinearity problems in estimation. Second, it is reasonable to assume that supercomputer buyers are primarily concerned with performance metrics and not so much with the machine specifics (e.g., number of processors).

Consumers can choose an "outside good", whose utility is $U_{r0t} = \xi_{0t} + \epsilon_{r0t}$. This outside alternative may represent not only other types of computers that buyers may consider powerful enough to meet their computing needs (e.g., mainframes, workstations), but also the flow utility from using a high-end computer one already owns. Therefore, it is reasonable to assume that the mean utility of the outside alternative, ξ_{0t}, evolves over time.

The following asusmption is imposed to obtain a closed-form solution for the market share of each good.

Assumption A5: $\forall j = 0, 1, .., J_t$, ϵ_{rjt} is identically and independently distributed with a Type I extreme-value distribution

Under assumption A5, the market share of product j at time t, $q_{jt}(p)$ is given by

$$q_{jt}(p) = \frac{\exp(\gamma x_j + \delta\kappa_{f(j)t} + \lambda MPP_j - \alpha p_{jt} + \xi_{jt} - \xi_{0t})}{1 + \sum_l \exp(\gamma x_l + \delta\kappa_{f(l)t} + \lambda MPP_l - \alpha p_{lt} + \xi_{lt} - \xi_{0t})}$$

and the outside good share is

$$q_{0t}(p) = \frac{1}{1 + \sum_l \exp(\gamma x_l + \delta\kappa_{f(l)t} + \lambda MPP_l - \alpha p_{lt} + \xi_{lt} - \xi_{0t})}$$

[11]Note that this state variable is bounded between 0 and 1. $\kappa_{f(j)t} = 0$ corresponds to the extreme case of inactivity in the industry, while $\kappa_{f(j)t} = 1$ indicates that firm f has produced the most powerful supercomputer up to time t.

3.1.2 Supply

Each firm f in the industry produces some subset $\mathcal{F}_{ft}$ of the $j = 1, ..., J_t$ supercomputer models commercialized at time t. Firm f 's profits from being active at time t are defined by

$$\Pi_{ft} = M_t \sum_{j \in \mathcal{F}_{ft}} (p_{jt} - mc_t(MPP_j, x_j))q_{jt}(p) - C_f$$

where C_f corresponds to fixed production costs. The marginal cost of producing supercomputer j at time t,is assumed to be time-varying and dependent on both observed quality x_j and whether j is an MPP system. To simplify notation, I denote mc_{jt} as the marginal cost function. M_t corresponds to the market size state at time t, which is defined as being the number of consumers potentially interested on purchasing a supercomputer. Firms are assumed to compete in prices in the spot market. I further assume the existence of an unique pure-strategy Nash-Bertrand equilibrium in prices[12]. Thus, in equilibrium, each price p_{jt} must satisfy the following first-order condition:

$$q_{jt}(p) + \sum_{k \in \mathcal{F}_{ft}} (p_{kt} - mc_{kt}) \frac{\partial q_{kt}(p)}{\partial p_{jt}} = 0$$

I use the fact that, under the multinomial logit demand derived above, the solution to the firm's first order conditions for each product j implies a constant absolute markup per firm. Anderson, de Palma and Thisse (1992, pp 251-252) prove that, for the multinomial logit demand, the unique solution to the system of firms' first-order conditions just derived satisfies the property $p_{jt} - mc_{jt} = p_{lt} - mc_{lt}$, $\forall j, l \in \mathcal{F}_{ft}$, $\forall f = 1, ..., F$. I build on the work of Nevo and Rossi (2006) and Rossi (2006) to derive both the average quality measure ν_{nt} and the profit function. Under the multinomial logit demand system, a buyer's expected utility from purchasing from firm f (henceforth *inclusive value of firm* f) is given by

$$\omega_{ft} = \ln \left(\sum_{l \in \mathcal{F}_{it}} \exp(\gamma x_l + \delta \kappa_{f(l)t} + \lambda MPP_l - \alpha p_{lt} + \xi_{lt} - \xi_{0t}) \right)$$

Inclusive values have been widely used as state variables in dynamic models since its development in McFadden (1981). Melnikov (2001) considers the inclusive value from buying any existing product for modelling consumer demand for printers. More recently, Hendel and Nevo (2006) used inclusive values from buying specific quantities to estimate demand for laundry detergent. However, this firm-specific quality measure is dependent on prices, which are assumed to be chosen by firms. Therefore, in order to use firm-specific inclusive values in my dynamic model, I must adjust for its endogeneity in prices. The constant

[12] Caplin and Nalebluff (1991) prove the existence of an unique equilibrium for this game under mild assumptions.

absolute markup property successfully accomplishes this objective. Let Mkp_{ft} be the constant absolute markup that firm charges for each of its supercomputers at time t. Using the fact that $\kappa_{f(j)t} = \kappa_{ft}$ for every set of products that firm f commercializes at time t, the firm inclusive value becomes

$$\omega_{ft} = \delta\kappa_{ft} - \alpha Mkp_{ft} + \ln(\nu_{ft})$$

where $\nu_{ft} \equiv \sum_{l \in \mathcal{F}_{it}} \exp(\gamma x_l + \lambda MPP_l - \alpha mc_{lt} + \xi_{lt} - \xi_{0t})$. Following Nevo and Rossi (2006) and Rossi (2006), the following result applies

Proposition 1 *The set of firms' first order conditions is equivalent to a set of F_t equations with F_t unknown equilibrium markups. This system implicitly defines the (unique) equilibrium markups as functions of the firm-specific observable states. Moreover, denoting $\nu_t \equiv \{\nu_{nt}\}_{n=1}^{N}$ and $\kappa_t \equiv \{\kappa_{nt}\}_{n=1}^{N}$,the equilibrium profit function of any firm f is given*

$$\Pi_{ft}(\mathbf{p}(\mathbf{s_t}), \mathbf{s}_t) = M_t \frac{\exp(-\alpha Mkp_{ft}(\nu_t, \kappa_t) + \ln(Mkp_{ft}(\nu_t, \kappa_t)) \exp(\delta\kappa_{ft})\nu_{ft}}{1 + \sum_{i=1}^{F_t} \exp(-\alpha Mkp_{it}(\nu_t, \kappa_t)) \exp(\delta\kappa_{it})\nu_{it}} - C_f$$

Proof: The existence of the unique equilibrium markups as functions of firm-specific observable states follows from Nevo and Rossi (2006). The proposed formula for equilibrium profits follows directly from replacing ω_{ft} by $\delta\kappa_{ft} - \alpha Mkp_{ft} + \ln(\nu_{ft})$ in the profit functions derivations of Nevo and Rossi (2006) and Rossi (2006).

3.2 Frontier benefits and innovation costs

3.2.1 Frontier benefits

As previously explained, supercomputer firms usually incorporate their advances in computing power into other lines of product (e.g. mainframes, PCs). Consequently, one would expect these benefits to be increasing on the firm's technological frontier. However, if there are several rivals with a similar technological state, the firm only gets a share of that benefit. For example, if two firms are competing on both the supercomputer and the PC market, the benefit of creating improved PCs by incorporating advances in supercomputing is lower if the rival is able to do the same. Therefore I assume that the benefit is decreasing in the number of technologically similar rivals. This motivates the following functional form for frontier benefits

$$\Upsilon_f(\mathbf{s}_t) = \frac{\rho_1 \kappa_{ft} + \rho_2 \kappa_{ft}^2}{\sum_{i=1}^{F} \mathbf{1}\left\{\kappa_{ft} - \phi \leq \kappa_{it} \leq \kappa_{ft} + \phi\right\}} \qquad \forall f = 1, .., F$$

where is a ϕ parameter defining how close is a rival i to a given firm f. Since κ_{ft} is a continuous variable I cannot restrict attention to the firms whose technological frontier is exactly the same as of firm f. Hence, the bandwidth parameter ϕ is required to define the number of technologically similar rivals. The variation in the number of a firm's similar rivals allows me to separate the benefit from innovating from the innovation costs. According to the "selection effect" introduced in the literature by Aghion et al. (2001), firm's incentives to innovate are increasing in the number of technologically similar rivals. As this number does not affect the costs of innovation, the benefit function can be identified from the data.

3.2.2 Innovation costs

I assume that, for all firms $f = 1, ..., F$, the functional form for innovation costs is given by:

$$C(I_{ft}, \mathbf{s}_t) = 1\{I_{ft} > 0\} \left(c_0 + \varepsilon_{1ft} + c_1 \kappa_{ft} + c_2 \sum_{i=1}^{F} \mathbf{1}\left\{\kappa_{it} \geq \kappa_{ft} + \phi\right\} (\kappa_{it} - \kappa_{ft}) + \right.$$

$$\left. c_3 \mathbf{1}\left\{\kappa_{ft} = 1\right\} + (c_4 + \varepsilon_{2ft}) I_{ft} + c_5 I_{ft}^2 \right)$$

The intuition behind this specification can be described as follows. Investment costs only take place if the firm extends its technological frontier. If the firm does expand its frontier, it pays a fixed cost which is assumed to depend on (i) a fixed amount c_0, (ii) a zero-mean shock ε_{1ft}, (iii) the firm's technological state κ_{ft}, (iv) the sum of the technological advantages of the firm's most advanced rivals, and (v) whether the firm is the technological leader. This specification aims to control firm heterogeneity and positive externalities on frontier expansion. The cost function is assumed to be quadratic in I_{ft}, in order to capture convexities in investment expenditures. The parameters c_3 and c_4 therefore provide information on marginal costs of innovation. Marginal costs are assumed to depend also on a private shock ε_{2ft}, but not dependent on firm heterogeneity. The average investment rate conditional on positive investment is approximatelly 300%, which suggests low marginal innovation costs.

3.3 State transitions

The specification of the expectations in the firm's problem requires assumptions on state transition functions. The evolution of M_t is determined by consumer behavior, while the realizations of each ν_{ft} depend on firm f's new product introduction policies. Modelling consumer behavior and new product introductions which do not expand the firm frontier is beyond the scope of this paper. Therefore, I assume that both market size and average quality states evolve stochastically under the following assumption:

Assumption A6: M_{t+1} is independent of ν_{ft+1}, for every $f = 1, ..., F$.

In order to guarantee that these assumptions are satisfied, parametric assumptions about these transitions are necessary. The fact that I have a relatively short time series (16 years) dictates a parsimonious specification for market size. Therefore I propose the following law of motion for M_t :

$$\ln(M_{t+1}) = \phi_0 + \phi_1 \ln(M_t) + \varsigma_{t+1}$$

where ς_{t+1} is assumed to be an independently and identically distributed error term. The regressor $\ln(M_t)$ controls intertemporal substitution. Consistent estimation can be performed using OLS, and the most suitable distribution for ς_{t+1} can be inferred from the residual properties.

Upon computing each ν_{it} by using the demand and supply estimates, its transition probability can be recovered after proper assumptions on its functional form. The evolution of this quality measure is determined by product introduction and scrappage decisions, whose modelling goes beyond the scope of this paper. Therefore, the law of motion for this state variable will be estimated nonparametrically using series estimators. In particular, I propose the following specification

$$\ln(\nu_{ft+1}) = \lambda_0 + \lambda_1 I_{ft} + \lambda_2 \sum_{j \neq f} I_{j,t} + \Psi_3(\boldsymbol{\lambda}_3; \ln(\nu_{ft})) + \Psi_4\left(\boldsymbol{\lambda}_4; \sum_{j \neq f} \nu_{j,t}\right) + \Psi_5\left(\boldsymbol{\lambda}_5; \ln(\nu_{nt}) \times \left(\sum_{j \neq n} \nu_{j,t}\right)\right)$$

$$+\Psi_6(\boldsymbol{\lambda}_6; M_t) + \Psi_7(\boldsymbol{\lambda}_7; \kappa_{ft}) + \Psi_8\left(\boldsymbol{\lambda}_8; \left(\sum_{j \neq n} \kappa_{j,t}\right)\right) + \Psi_9\left(\boldsymbol{\lambda}_9; \kappa_{ft} \times \left(\sum_{j \neq n} \kappa_{j,t}\right)\right)$$

$$+\Psi_{10}\left(\boldsymbol{\lambda}_{10}; \kappa_{ft} \times \left(\sum_{j \neq n} \nu_{j,t}\right)\right) + \varrho_{f,t+1}$$

where $\Psi_i(\boldsymbol{\lambda}_i, y)$ denotes a series of polynomials on variable y with coefficient vector $\boldsymbol{\lambda}_i$. $\varrho_{f,t+1}$ is assumed to be an i.i.d. zero-mean error. The resulting equation can be estimated using OLS, and the residual properties of $\varrho_{f,t+1}$ can be used to propose a distribution function for this error term.

I am left to define the transition of a state that describes the evolution of the firms' technological frontier. Conditional on states and actions, the transition for is deterministic. The following result establishes its transition function

Proposition 2 *The transition of κ_{ft+1} conditional on player's actions and observed states is deterministic and given by*

$$\kappa_{ft+1} \equiv \frac{(1+I_{ft})\kappa_{ft}}{\max\limits_{i=1,\dots,F}\{(1+I_{it})\kappa_{it}\}}$$

Proof: Let h_{ft} be firm f's frontier at time t. Next period's frontier, h_{ft+1}, is simply $h_{ft}(1+I_{ft})$. This establishes the numerator of the transition. The industry's frontier in the next period if simply the maximal h_{it+1} after investment decisions. By definition, we have

$$\kappa_{ft+1} \equiv \frac{h_{ft+1}}{\max\limits_{i=1,\dots,F}\{(h_{it+1}\}} = \frac{(1+I_{ft})h_{ft}}{\max\limits_{i=1,\dots,F}\{(1+I_{it})h_{it}\}}$$

The transition of κ_{ft+1} is the obtained by dividing both the numerator and the denominator by $\max\limits_{i=1,\dots,F}\{h_{it}\}$.

3.4 Equilibrium concept

The logit structure of demand allows me to assume that firms choose constant absolute markups rather than specific retail prices. Therefore, an equilibrium concept can be defined by considering Markovian strategies of the form $\sigma_f(\mathbf{s},\varepsilon_f) \equiv \{Mkp_f(\mathbf{s}), I_f(\mathbf{s},\varepsilon_f)\}$. Let $V_f(\mathbf{s}|\sigma_f,\sigma_{-f})$ be firm f's expected discounted payoffs when he chooses strategy σ_f and his rivals choose σ_{-f}. A Markov-Perfect Nash Equilibrium (MPNE) of this dynamic oligopoly game can be defined as follows

Definition 1 (MPNE): A Markovian strategy profile $\sigma^* \equiv (\sigma_1^*, ..., \sigma_F^*)$ is an MPNE if, for every firm f, σ_f^* solves f's problem given σ_{-f}^*. That is, $V_f(\mathbf{s}|\sigma_f^*,\sigma_{-f}^*) \geq V_f(\mathbf{s}|\hat{\sigma}_f,\sigma_{-f}^*)$ for all $\mathbf{s}$ and alternative strategy $\hat{\sigma}_f$.

The existence of MPNE follows from Proposition 1 in Doraszelski and Satterthwaite (2005). In order for the MPNE to be of practical interest, it is necessary to check its computational feasibility and uniqueness. In general, solving for the MPNE without additional conditions on the equilibrium is computationally prohibitive. This motivates the introduction of the following concepts:

Definition 2 (Symmetry): A function f_n is said to be symmetric if $f_n(z_1, ..., z_n, ..., z_N) = f_1(z_n, ..., z_1, ..., z_N)$ for all $n = 1, ..., N$

Definition 3 (Anonymity): A function f_n is said to be anonymous if $f_1(z_1, z_2..., z_k, ..., z_l, ..., z_N) = f_1(z_1, z_2..., z_l, ..., z_k, ..., z_N)$ for all $k \geq 2$ and $l \geq 2$.

The following definition of symmetric and anonymous MPNE is based on the framework of Doraszelski and Satterthwaite (2005):

Definition 4 (Symmetric and Anonymous MPNE): An MPNE is said to be symmetric and anonymous if the both equilibrium policies and value functions are symmetric and anonymous for all F players.

The computational burden of solving an MPNE is considerably reduced if it is symmetric and anonymous. Fortunately, Doraszelski and Satterthwaite (2005) prove the existence of a symmetric and anonymous MPNE.

Corollary 1 *There is a symmetric and anonymous MPNE for the game.*

Proof of corollary 1: If follows trivially that the parametric forms for profits, benefits, and innovation costs satisfy anonymity and symmetry. The assumptions on the parametric forms for the transition probabilities of observable states satisfy anonymity. The result is a direct application of Lemma 2 in Doraszelski and Satterthwaite (2005).

For the rest of the paper, I assume that there is an unique symmetric and anonymous MPNE[13].

4 Estimation methods

The structural estimation of the model consists of a variant of the two-step method proposed by Bajari, Benkard and Levin (2006). I start by providing a general overview of my estimation strategy, explaining its differences with respect to the method of Bajari et al. Then I discuss the technical details of each step of the estimation procedure.

[13] The conditions for uniqueness of symmetric and anonymous MPNE are still an open area of research in this literature. Since optimal entry and exit policies are unique, one could attempt to prove uniqueness of MPNE for this dynamic game by checking if there is an unique solution for the firm Kuhn-Tucker conditions on investment. A similar approach is followed by Seim (2001). However, this strategy of proof generally requires conditions on the game parameters, which are unknown to the researcher.

4.1 An overview of the estimation

The first step of my estimation strategy coincides with the one of Bajari, Benkard and Levin (2006). It consists of estimating the profit function parameters, as well as optimal policies and transition functions. After choosing appropriate instruments, the profit function parameters are recovered by estimating the supply and demand systems. The estimates are then used to compute the average quality state ν_{ft} for all players $f = 1, ..., F$. Optimal policies and transition probabilities for M_t and each ν_{ft} are nonparametrically estimated using polynomial series of observed states.

In the second step of estimation, both my approach and the one of Bajari et al. use the linearity of the firm's value function on the parameters not recovered in the first stage. This property allows me to decompose that value function into the vector of remaining parameters and the vector of expected discounted payoffs and actions. Therefore, the entries of the latter need to be computed only once by forward simulation of payoffs and actions using the estimated policies and transitions. However, the two approaches differ from this point on. Bajari et al. estimate the vector of remaining parameters by using the definition of MPNE outlined in the previous section. After generating a set of I alternative policies, $\left\{\hat{\sigma}_f^{(i)}\right\}_{i=1}^{I}$, the vector of parameters is estimated by minimizing the number of violations of the condition $V_f(\mathbf{s}|\sigma_f^*, \sigma_{-f}^*) \geq V_f(\mathbf{s}|\hat{\sigma}_f^{(i)}, \sigma_{-f}^*)$.

The variant of this method followed in this paper instead uses the linearity of both the value function and its derivatives to replace unknown components in firm optimality conditions on investment. The optimal investment policy is fully defined by two conditions: (i) marginal benefit of innovation must equal its marginal cost, and (ii) investment only takes place if it is better than zero investment. Optimal investment is monotonic in the private information shocks[14], implying that the investment distribution conditional on observed states is identified from the data. This fact allows me to estimate the remaining parameters of the model with maximum-likelihood estimation. In the spirit of the Tobit I model, I separate the case where firms make zero investments from the case where it is strictly positive to form a likelihood function. As shown below, the first-order condition on investment depends on both the value function and its derivative with respect to the firm's investment rate. These are linear in all the remaining parameters of the model. Therefore, I am able to decompose both functions into the vector of remaining parameters and the vector of expected discounted payoffs and actions. This approach has the advantage of avoiding three sources of inneficiency in the estimation method proposed by Bajari et al.: (i) a suboptimal choice of the inequality sampling distribution, (ii) sampling error from considering a finite number of sampled inequalities.

[14]This result follows from the proposed functional form for innovation costs. The fact that optimal investment is decreasing in the shock on marginal costs follows from the Implicit Function Theorem applied to the first-order condition on investment (see section 4.3 and Appendix B). Increases in the shock on fixed innovation costs imply either no-investment, or the investment level implied by the first-order condition.

4.2 First-step estimates

4.2.1 Profit function parameters

Given the logit structure of demand, I follow the procedure proposed by Berry (1994) to obtain consistent estimates of its parameters. I assume that the unobserved supercomputer quality can be modelled as $\xi_{jt} = \xi_j + \Delta\xi_{jt}$. That is, there are unobserved attributes of the supercomputer that are constant over time (e.g., design, power consumption) and perceived quality features that might vary by market[15] (e.g., changes in machine reliability, news about the supercomputer ability to solve problems in selected applications). Under this assumption, the market share for product j at time t becomes

$$q_{jt}(p) = \frac{\exp(\gamma x_j + \delta\tilde{\kappa}_{f(j)t} + \lambda MPP_j - \alpha p_{jt} + \xi_j - \xi_{0t} + \Delta\xi_{jt})}{1 + \sum_l \exp(\gamma x_l + \delta\tilde{\kappa}_{f(l)t} + \lambda MPP_l - \alpha p_{lt} + \xi_l - \xi_{0t} + \Delta\xi_{lt})}$$

By taking the natural logarithm of the ratio between product j's share and the outside good share, the following estimation equation is obtained

$$\ln(q_{jt}) - \ln(q_{0t}) = \gamma x_j + \delta\tilde{\kappa}_{f(j)t} + \lambda MPP_j - \alpha p_{jt} + \xi_j - \xi_{0t} + \Delta\xi_{jt}$$

For the empirical application, ξ_{0t} can be controlled in estimation by using time fixed effects. One can estimate ξ_j by assuming either random or fixed effects. In this paper, I follow the random effects specification. The term $\Delta\xi_{jt}$, which corresponds to unobserved quality features, can be treated as econometric error term. However, this term should be correlated with prices, since manufacturers take into account all product characteristics in their pricing decisions. In addition, there is also reason to believe that the observed quality measure, x_j, is correlated with the error term. This is because unobserved characteristics may be contributing to the supercomputer's ability to solve a system of 1000 equations with 1000 unknowns, which is measured by Rmax. Therefore, valid instruments for both prices and observed quality should are necessary. Determining how relevant are these endogeneity issues for estimation is an empirical issue. Therefore, the details about the choice of instruments will be left for the empirical results section.

Marginal cost parameters can only be recovered from the supply side of the product market by using the firms' first order conditions on prices. To see this, note that

$$\frac{\partial q_{kt}(p)}{\partial p_{jt}} = \begin{cases} \alpha q_{jt}(p)[q_{jt}(p) - 1] & \text{if } k = j \\ \alpha q_{jt}(p) q_{kt}(p) & \text{if } k \neq j \end{cases}$$

Again denoting $Mkp_{ft} = p_{jt} - mc_{jt}$, $\forall j \in \mathcal{F}_{ft}$, each firm's first order condition with respect to p_{jt} becomes

[15] Recall that each period t is considered a different market.

$$Mkp_{ft} = \frac{1}{\alpha}\frac{1}{1-q_{ft}(p)} \implies p_{jt} = \frac{1}{\alpha}\frac{1}{1-q_{ft}(p)} + mc_{jt}$$

where $q_{ft}(p) = \sum_{k \in \mathcal{F}_{ft}} q_{kt}(p)$, i.e., the total market share of firm f. In order to estimate marginal costs, assumptions on its functional form are necessary. I assume that marginal costs are a linear function of observed quality, x_j, architecture type (i.e., whether the supercomputer is an MPP) and of a set of time-varying exogenous parameters, η_t. That is, $mc_{jt} = \eta_0 + \eta_t x_j + \eta_{MPP} MPP_j + \vartheta_{jt}$, where ϑ_{jt} is a shock to marginal costs assumed independent and identically distributed over time and products.The inclusion of η_t is justified by reported evidence on supercomputer imput prices decline, namely from the increasing average number of processors and processor speed. There are two possible ways to recover marginal cost parameters: (i) replacing the value of α from demand estimation in the last markup equation, and use OLS after moving the term in $\frac{1}{\alpha}$ to the left-hand side of the markup equation, and (ii) jointly estimate the demand and supply equations using GMM with proper instruments. I followed approach (i).

4.2.2 Investment Policy

Since optimal absolute markups can be recovered directly from demand and marginal cost estimates by solving the system of F asbolute markups, only the optimal policy for the investment rate needs to be recovered from the data. I estimate a reduced form for the investment policy using nonparametric methods. However, I must control for the lumpy investment behavior seen in the data. Even though firms decide optimal investment rates by comparing marginal benefits to marginal costs of innovation, they may not invest due to considerable fixed costs of innovation. In addition, my estimation strategy must account for the fact that firms cannot make negative investments. This motivates the estimation of the investment policy by using the Tobit I model. Like for the average quality transition estimation, I considered polynomials on observed states as regressors. Since the number of technologically similar rivals influences the benefit of innovation, I also considered polynomial series in this variable as additional regressors. The proposed specification for the investment policy is therefore $I_{ft} = \max\{0, w_{f1}(\mathbf{s}_t, \boldsymbol{\theta}) + \eta_{ft}\}$, where η_{ft} is a zero-mean normally distributed error and

$$w_{f1}(\mathbf{s}_t, \boldsymbol{\theta}) = \theta_0 + \Psi_{11}(\boldsymbol{\theta}_1; \ln(\nu_{ft})) + \Psi_{12}\left(\boldsymbol{\theta}_2; \sum_{j \neq f} \ln(\nu_{j,t})\right) + \Psi_{13}\left(\boldsymbol{\theta}_3; \ln(\nu_{nt}) \times \left(\sum_{j \neq n} \ln(\nu_{j,t})\right)\right)$$

$$+\Psi_{14}(\boldsymbol{\theta}_4; M_t) + \Psi_{15}(\boldsymbol{\theta}_5; \kappa_{ft}) + \Psi_{16}\left(\boldsymbol{\theta}_6; \left(\sum_{j \neq n} \kappa_{j,t}\right)\right) + \Psi_{17}\left(\boldsymbol{\theta}_7; \kappa_{ft} \times \left(\sum_{j \neq n} \kappa_{j,t}\right)\right)$$

$$+\Psi_{18}\left(\boldsymbol{\theta}_8;\kappa_{ft}\times\sum_{j\neq f}\ln(\nu_{j,t})\right)+\Psi_{19}\left(\boldsymbol{\theta}_9;\ln\left(\sum_{i=1}^{F}\mathbf{1}\left\{\kappa_{ft}-\phi\leq\kappa_{it}\leq\kappa_{ft}+\phi\right\}\right)\right)$$

The choice of the bandwidth parameter ϕ is an empirical issue, and hence left for the estimation results section.

4.3 Second-step estimates

The first-step recovers all parameters except the ones of the investment cost function and frontier positioning benefit. The first-order conditions on optimal investment provide essentially information of marginal investment costs and benefits. In order to obtain more information on the fixed investment costs, one must impose the condition that if a positive investment rate is chosen, then it must be better than no investment. For ease of exposition, the details on both the derivation of these conditions and the value function decomposition can be found in Appendix B. I proceed by deriving the log-likelihood function used in estimation. The first-order condition on investment is given by

$$I_f\left[-\frac{\partial C(I_f,\mathbf{s},\varepsilon_f)}{\partial I_f}+\beta\int_{\varepsilon_{-f}}E_{\mathbf{s},,\sigma_f,\sigma_{-f}}\left(\frac{\partial V_f(\mathbf{s}',\varepsilon'_f)}{\partial\kappa'_f}\frac{\partial\kappa'_f}{\partial I_f}+\sum_{i\neq f}\frac{\partial V_f(\mathbf{s}',\varepsilon'_f)}{\partial\kappa'_i}\frac{\partial\kappa'_i}{\partial I_f}\right)dG(\varepsilon_{-f})\right.$$
$$\left.+\beta\int_{\varepsilon_{-f}}E_{\mathbf{s},,\sigma_f,\sigma_{-f}}\left(V_f(\mathbf{s}',\varepsilon'_f)\frac{\partial\ln\left(f_{\nu'}(\nu'|\mathbf{s},I_f,\sigma_{-f}(\mathbf{s},\varepsilon_{-f}))\right)}{\partial I_f}\right)dG(\varepsilon_{-f})\right]=0$$

Let θ denote the vector of remaining parameters, and let $\tilde{\theta}=[1,\theta]$. In addition, let $R(\mathbf{s}',I_f,\varepsilon'_f)$ and $W(\mathbf{s}',I_f,\varepsilon'_f)$ be two vectors of discounted terms so that

$$R(\mathbf{s}',I_f,\varepsilon'_f)\ \tilde{\theta}=\frac{\partial V_f(\mathbf{s}',\varepsilon'_f)}{\partial\kappa'_f}\frac{\partial\kappa'_f}{\partial I_f}+\sum_{i\neq f}\frac{\partial V_f(\mathbf{s}',\varepsilon'_f)}{\partial\kappa'_i}\frac{\partial\kappa'_i}{\partial I_f}$$

$$W(\mathbf{s}',I_f,\varepsilon'_f)\tilde{\theta}=V_f(\mathbf{s}',\varepsilon'_f)\frac{\partial\ln\left(f_{\nu'}(\nu'|\mathbf{s},I_f,\sigma_{-f}(\mathbf{s},\varepsilon_{-f}))\right)}{\partial I_f}$$

then the firs-order condition on investment can be rewritten as

$$I_f\left[-\frac{\partial C(I_f,\mathbf{s},\varepsilon_f)}{\partial I_f}+\beta(\hat{R}(\mathbf{s},I_f)+\hat{W}(\mathbf{s},I_f))\tilde{\theta}\right]=0$$

where $\hat{R}(\mathbf{s},I_f)=\int_{\varepsilon_{-f}}E_{\mathbf{s},,\sigma_f,\sigma_{-f}}\left(R(\mathbf{s}',I_f,\varepsilon'_f)\right)dG(\varepsilon_{-f})$ and $\hat{W}(\mathbf{s},I_f)=\int_{\varepsilon_{-f}}E_{\mathbf{s},,\sigma_f,\sigma_{-f}}\left(W(\mathbf{s}',I_f,\varepsilon'_f)\right)dG(\varepsilon_{-f})$
. The term inside square brackets implicitly defines optimal investment conditional on an interior nonnegative solution. If the firm decides to innovate, then it must be the case that the discounted payoff under optimal investment is greater than not investing at all. This corresponds to the following condition

$C(I_f^*, \mathbf{s}, \varepsilon_f) + \beta \hat{Y}_\sigma(\mathbf{s})\tilde{\theta} > \beta \hat{Y}_{\sigma_{-f} I=0}(\mathbf{s}))\tilde{\theta}$

where I use the fact that the value function can be linearly decomposed as $V_f(\mathbf{s}', \varepsilon_f') = Y_\sigma(\mathbf{s}', \varepsilon_f')\tilde{\theta}$, and $\hat{Y}_\sigma(\mathbf{s}) = \int\limits_{\varepsilon_{-f}} E_{\mathbf{s},,\sigma_f,\sigma_{-f}} \left(Y_\sigma(\mathbf{s}', \varepsilon_f')\tilde{\theta} \right) dG(\varepsilon_{-f})$. The discounted value in case of no investment, $\hat{Y}_{\sigma_{-f} I=0}(\mathbf{s}))$, is analogously defined. Note that this condition can be used to define the probability of no-investment. In addition, note that the term in square brackets in the first-order condition implicitly defines the density of optimal investment conditional on the latter being positive. Therefore, I can construct a joint likelihood function of the form

$$L = \prod_{t=1}^{T} \prod_{f=1}^{F} \left[(\Pr(I_{ft} = 0|\mathbf{s}_t))^{(1-y_{1ft})} f(I_{ft}|\mathbf{s}_t)^{y_{1ft}} \right]$$

where y_{1ft} is a dummy variable which equals unity if firm f makes strictly positive innovation at time t and equals zero otherwise. So the objective is to define $\Pr(I_{ft} = 0|\mathbf{s}_t)$ and $f(I_{ft}|\mathbf{s}_t)$ in terms of θ and estimate this parameter vector by maximizing the logarithm of function L. By using the functional form assumption on innovation costs, the first-order condition on investment for interior solutions simplifies to

$$c_4 + \varepsilon_{2ft} + 2c_5 I_{ft} + \beta(\hat{R}(\mathbf{s}, I_f) + \hat{W}(\mathbf{s}, I_f))\tilde{\theta} = 0$$

So the distribution of optimal investment is implicitly defined by the distribution of ε_{2ft}. For the case of the "investment-no investment" condition, we have

$$c_0 + \varepsilon_{1ft} + c_1 \kappa_{ft} + c_2 \sum_{i=1}^{F} \mathbf{1}\{\kappa_{it} \geq \kappa_{ft} + \phi\} (\kappa_{it} - \kappa_{ft})$$

$$+ c_3 \mathbf{1}\{\kappa_{ft} = 1\} + (c_4 + \varepsilon_{2ft}) I_{ft} + c_5 I_{ft}^2 + \beta(\hat{Y}_\sigma(\mathbf{s}) - \hat{Y}_{\sigma_{-f} I=0}(\mathbf{s}))\tilde{\theta} > 0$$

However, this condition is only useful for the case where observed investment is zero. Since I_{ft} in this equation regards optimal investment conditional on an interior choice, this condition must be corrected for two problems: (i) the investment conditional on interior choice is only observed if positive investment occurs in the data, and (ii) this investment rate depends on ε_{2ft}. I proceed with this correction by replacing I_{ft} with the first-stage estimate of the investment policy. The tobit structure of the estimated investment policy allows me to compute $E[I_{ft}|I_{ft} > 0, \mathbf{s}_t]$. Hence, under suitable assumptions on the densities for ε_{1ft} and ε_{2ft}, the probability of positive investment can be computed. Under the assumption that for ε_{1ft} and ε_{2ft} are indenpendently and identically distributed zero-mean normal with standard deviations σ_1 and σ_2, we have

$$\Pr(I_{ft} > 0|\mathbf{s}_t) = \int_{\varepsilon_{2ft}-\infty}^{\infty} \Phi\left(\frac{\Theta\left(\theta, \varepsilon_{2ft}, E[I_{ft}|I_{ft} > 0, \mathbf{s}_t]\right)}{\sigma_1}\right) \phi(\varepsilon_{2ft}, \sigma_2) d\varepsilon_{2ft}$$

where $\Theta\left(\theta, \varepsilon_{2ft}, E[I_{ft}|I_{ft} > 0, \mathbf{s}_t]\right)$ corresponds to the LHS of the "investment-no investment" condition with I_{ft} replaced by $E[I_{ft}|I_{ft} > 0, \mathbf{s}_t]$. Therefore, the remaining parameters of the model can be estimated by maximizing the following log-likelihood function

$$\theta_{MLE} = \arg\max_{\theta} \sum_{t=1}^{T} \sum_{f=1}^{F} \left[y_{1ft} \ln\left(\frac{1}{\sigma_2}\phi\left(\frac{-c_4 - 2c_5 I_{ft} - \beta(\hat{R}(\mathbf{s}, I_f) + \hat{W}(\mathbf{s}, I_f))\tilde{\theta}}{\sigma_2}\right)\right) + (1 - y_{1ft}) \ln\left(1 - \Pr(I > 0)\right)\right]$$

The details on simulation methods required to compute the discounted function terms can be found in Appendix B.

5 Data and estimation results

5.1 The data

The data required for estimation using the methods described in the previous section consists of the following variables: market shares and prices in each year, supercomputer characteristics, market size (M_t), technological rankings of firms (i.e., κ_{nt}), innovation rates. The data ranges from 1990 to 2005. The details on the data construction and the definition of its variables are presented in Appendix A.

Tables II and III provide yearly descriptive statistics of the most relevant variables in the data set. Except for the first four columns (measured in units), all entries on Table II correspond to sales-weighted averages, where real prices are in \$1M units, Rmax is measured in GFLOPS, and processors speed is measured in MHz. Firmrank κ_{nt} corresponds to the ratio between a firm's maximal Rmax production record and the industry's maximum. Some relevant trends can be detected. First, there is a considerable growth in total sales up to 1997, followed by a stable evolution of the quantity sold. This pattern follows closely the evolution of the average real price statistic, which experiences a sharp decline up to the mid 1990s and then varies between \$6M and \$10M. There is a exponential growth on the average observed quality of supercomputers, and a similar evolution applies to the average number of processors and processor speed used in the machines. Finally, the wave-like evolution of the average firmrank indicates that firms tend to quickly approach the market leader in quality record until some industry player performs a considerable innovation. The descriptive statistics presented on Table III not only reinforce this evidence, but also indicate that firms tend to innovate considerably, even though no innovation is particularly frequent.

TABLE II

YEARLY SALES-WEIGHTED MEANS

Year	Quantity	Market size	Price	Rmax	No. Processors	Proc. speed	Firmrank (κ_{nt})
1990	73	134	11.549	1.578	215.246	140.058	0.129
1991	101	234	7.793	1.733	151.435	184.961	0.230
1992	169	399	8.455	2.508	51.443	137.896	0.428
1993	275	670	7.363	5.212	189.646	112.010	0.281
1994	400	1059	4.708	6.249	87.955	108.876	0.390
1995	405	1444	2.706	6.590	43.049	107.524	0.398
1996	407	1830	6.245	19.391	61.985	180.838	0.322
1997	584	2341	9.581	42.021	111.603	246.689	0.177
1998	407	2647	8.457	51.801	130.081	275.771	0.378
1999	510	2988	9.471	94.920	192.558	339.015	0.401
2000	477	3190	8.789	145.313	184.972	389.066	0.478
2001	408	3198	7.333	242.820	227.902	536.621	0.387
2002	449	3242	8.968	638.469	347.791	1006.972	0.126
2003	388	3223	6.865	836.986	341.118	1650.017	0.229
2004	489	3128	9.552	1889.182	614.918	1951.751	0.473
2005	463	3184	9.961	4329.1	1382.551	2529.151	0.417

TABLE III

SUMMARY STATISTICS

Variable	Mean	Median	Std	Min	Max
Price	7.952	5.660	9.791	0.238	317.126
Rmax	633.465	51.2	4683.404	0.422	280600
No. processors	298.336	96	2092.405	1	131072
Processor speed	713.410	333.3	893.827	7	3600
Product market share(%)	0.2	0.04	0.4	0.031	5.3
Firm market share(%)	4.1	4.5	3.1	0.031	26.1
Innovation rate (%)	296.4	47.8	857.2	0	8300.7
Firmrank (κ_{ft})	0.424	0.240	0.368	.006	1
Market size (M_t)	2056.935	2647	1215.986	134	3242

5.2 Estimation Results

I estimate demand and supply equations separately. In order to assess the need for instruments in demand estimation, I run a preliminary OLS regression of the logit demand equation. In the first column of Table IV, I report the results of this regression. The coefficients on product characteristics are unintuitive. For example, one would expect marginal utility to be increasing in observed quality. This suggests that this variable is correlated with unobserved attributes. Moreover, the coefficient on price has the expected sign, but implies that all supercomputer models have inelastic demands. This contradicts profit maximizing behavior. Therefore, the OLS results indicate that instrumental variables are required for consistent estimation of demand.

In addition to the problem of price endogeneity, I must account for the endogeneity of observed quality (Rmax) suggested by the OLS regression. Therefore, instruments must be carefully chosen in order to avoid correlation with unobserved attributes. One alternative is to assume that product characteristics are predetermined and therefore suitable instruments. Even though the richness of the Top500 data would allow me to consider the attributes ignored in the utility specification as instruments, only three of those were considered: processor speed, a dummy on whether the processor was produced by Cray, and a dummy variable for cluster systems (i.e. supercomputers which result of clustering less powerful systems).

The speed of processors included in supercomputers is primarily explained by the exogenous evolution in processor technology during the 1990s. As processors are a key input of supercomputers, processor speed should be correlated with both price and observed quality. Moreover, the fact that some manufacturers use processors produced within the firm for some of their supercomputers (e.g. Cray) should be influencing the retail price. However, a dummy variable for Cray processors is unlikely to be correlated with the supercomputer unobserved characteristics, since about 7% of the machines in the sample use Cray processors[16]. The clustering property is shared by about 15% of the supercomputers in the sample, so it is unlikely to be correlated with the machine-specific unobserved attributes. Most of other variables in the data (e.g., number of processors) are not likely to be valid instruments. For example, supercomputers with more processors may be preferred for applications where it is convenient to assign separate parts of the problem being solved to each processor. Treating (some) characteristics of the product as exogenous is as reasonable here as in previous work (e.g., Berry, Levinsohn and Pakes(1995), Nevo (2001)).

[16] As a robustness check, I runned the IV logit regression without the dummy for Cray processors. The effect in the final estimates is negligible. However, this dummy variable is significant in the first-stage regressions. Therefore I preferred the instrument set where this dummy is included.

TABLE IV

LOGIT DEMAND ESTIMATES

	OLS Logit Demand		IV Logit Demand	
Variable	Estimate	Std. Error	Estimate	Std. Error
Price (in $1M)	-0.003	0.002	–0.374	0.193
Firmrank	-0.229	0.052	1.949	1.069
Rmax (in Teraflops)	-0.012	0.028	2.462	0.139
MPP dummy	-0.244	0.037	0.488	0.423
Constant	-3.298	0.142	-1.643	2.594
1991	-0.804	0.186	-1.793	0.588
1992	-1.360	0.171	-3.195	0.958
1993	-1.879	0.161	-4.208	1.230
1994	-2.308	0.156	-5.771	1.715
1995	-2.743	0.158	-6.465	1.865
1996	-3.217	0.153	-6.214	1.484
1997	-3.465	0.151	-5.338	0.985
1998	-3.733	0.153	-6.456	1.427
1999	-3.769	0.152	-7.443	1.884
2000	-3.858	0.153	-7.849	2.059
2001	-3.870	0.156	-8.697	2.393
2002	-3.917	0.155	-7.952	2.222
2003	-4.155	0.154	-9.124	2.763
2004	-3.903	0.157	-10.882	3.908
2005	-3.942	0.167	-12.849	5.309
Adj. R^2	0.592		n.a.	
% inelastic demands	100%		11.23%	
Number of obs.	2204		2204	

In order to ensure that on endogeneity concerns arise, it would be convenient to consider alternative sources of instruments. Cost shifters are an obvious candidate. McCallum (2002) provides series for memory prices for the period 1957-2006 at www.jcmit.com. Based on this series, I compute the average price of a Megabyte of memory (in 1998 dollars) for 1990-2005, and use it in the instrument set along with cluster dummy and processor speed. Demand estimation was performs using panel 2SLS regressions with random effects. The unit for prices is \$1M, in real 1998 dollar units, while the measure for Rmax is Teraflops (1000 Gigaflops). Results are presented on table IV. All the parameters have the expected sign and magnitude. The significance and negative signs of the time effects is in line with the rapid obsolescence of supercomputers due to rapid innovation in the industry and incorporation of technical advances in other goods (e.g. mainframes, workstations).

For the supply equation estimates, I assumed that $mc_{jt} = \psi_t x_{jt} + \psi MPP_j + \zeta_{jt}$, where η_t corresponds to a vector of time dummies multiplying observed quality and ζ_{jt} is a noise term. Results of marginal cost estimation from the supply equation using OLS -both with and without the dummy for MPP systems - are presented on Table V[17]. The coefficients corresponds to \$1M in 1998 dollars. From the results, one can infer that marginal costs for producing a Teraflop are decreasing over time under both specifications. This evidence agrees with two facts about supercomputer imputs: (i) processors are increasingly cheaper and of improved speed, and (ii) memory prices per Megabyte are falling over time[18], even though the improvements in memory technology are far less considerable than the ones for processors. The MPP dummy essentially affects the constant part of the marginal costs, but adds very little explanatory power. However, since its coefficient is statistically significant, the specification with MPP dummy was preferred.

Another important implication to be derived from the marginal cost estimates concerns its relation with Moore's Law. Even though Moore's original statement concerned transistor technology, it is common to cite Moore's Law to refer to the rapidly continuing advance in computing power per unit cost. In the context of computing technology, Moore's Law prescribes that the computing power per unit cost doubles every 18 months[19]. With the estimates of marginal cost at hand, it is possible to check this law by estimating the temporal evolution of marginal costs for a fixed level of computing power. The results for marginal costs suggest a decreasing exponential evolution over time, I fit the following equation by OLS for the data for both columns of Table V

$$\ln(mc_t) = \mu_0 + \mu_1 t + v_t \quad t = 1990, ..., 2005$$

Table VI presents the results for the temporal evolution marginal costs of producing 1 Teraflop. The results for the two sets of estimates are very similar, so I will focus the analysis on the results of the second column.

---

[17]The coefficient multiplying the inverse share of the rivals, $\frac{1}{\alpha}$, was replaced in the markup equation using the demand estimate for α. A GMM estimation strategy where both demand and supply equations are jointly considered (e.g., Berry, Levinsohn and Pakes(1995)) can be followed in future developments.

[18]Source: http://www.jcmit.com/mem2006.htm

[19]This is the most popular version of Moore's Law, even though Moore's original claim was that transistor density of integrated circuits, with respect to minimum component cost, doubles every 24 months.

TABLE V

MARGINAL COST ESTIMATES

	without MPP dummy		with MPP dummy	
Variable	Estimate	Std. Error	Estimate	Std. Error
Rmax×1990	2264.987	237.816	2296.836	237.229
Rmax×1991	1266.243	231.524	1311.899	231.119
Rmax×1992	1706.449	194.590	1715.699	194.005
Rmax×1993	325.681	69.697	315.194	69.536
Rmax×1994	251.732	44.743	236.372	44.786
Rmax×1995	211.773	36.236	206.399	36.152
Rmax×1996	261.401	9.235	258.551	9.237
Rmax×1997	200.479	4.230	198.155	4.261
Rmax×1998	124.753	5.298	122.740	5.308
Rmax×1999	60.201	3.073	58.720	3.088
Rmax×2000	29.650	2.004	28.548	2.018
Rmax×2001	17.913	1.482	17.362	1.484
Rmax×2002	8.819	0.750	8.585	0.750
Rmax×2003	4.172	0.554	4.328	0.554
Rmax×2004	2.463	0.321	2.629	0.323
Rmax×2005	1.392	0.150	1.424	0.150
Constant	1.283	0.169	0.827	0.206
MPP	-	-	0.992	0.260
Adjusted R^2	0.6325		0.6348	

TABLE VI

MARGINAL COSTS AND MOORE'S LAW

	without MPP dummy		with MPP dummy	
Variable	Estimate	Std. Error	Estimate	Std. Error
Constant	897.102	44.156	952.900	50.477
Year	-0.447	0.022	-0.475	0.025
Adjusted R^2	0.9645		0.959	

The instantaneous rate of decrease in marginal costs is 47.5%, which implies an anual decrease in marginal costs of approximately 38.7%. So the ratio of computing power over unit cost is increasing by 63.132% per year. Therefore, the marginal cost estimates closely resemble Moore's Law.

TABLE VII

INVESTMENT POLICY ESTIMATES

Variable	Estimates	Std. Error	Estimates	Std. Error	Estimates	Std. Error
Constant	-5.294	12.995	-1.751	12.451	0.277	12.064
log-average quality ($\ln(\nu_{ft})$)	-0.401	0.166	-0.389	0.167	-0.370	0.165
Rival log-average quality $\left(\sum_{i \neq f} \ln \upsilon_{it}\right)$	-0.076	0.030	-0.074	0.030	-0.073	0.030
$\ln(\nu_{ft}) \times \left(\sum_{i \neq f} \ln \upsilon_{it}\right)$	-0.009	0 .004	-0.009	0 .004	-0.009	0 .004
Firm rank (κ_{ft})	-71.493	36.025	-78.926	35.353	-85.218	34.485
Squared firm rank (κ_{ft}^2)	34.443	22.590	37.901	22.372	41.350	21.928
Rivals' aggregate rank $\left(\sum_{i \neq f} \kappa_{it}\right)$	7.257	5.111	6.713	5.099	6.744	5.106
Squared rivals' aggregate rank	-1.369	0.750	-1.304	0.751	-1.337	0.753
rank cross product $(\kappa_{ft}) \times \left(\sum_{i \neq f} \kappa_{it}\right)$	14.895	4.936	15.763	4.899	16.516	4.884
Squared rank cross product	-5.703	1.729	-5.902	1.732	-6.172	1.737
Cube of rank cross product	1.389	0.668	1.515	0.656	1.787	0.661
$(\kappa_{ft}) \times \left(\sum_{i \neq f} \ln \upsilon_{it}\right)$	3.569	1.436	3.576	1.453	3.620	1.455
Squared rank and rival quality product	3.303	0.946	3.233	0 .947	3.221	0.952
Cube of rank and rival quality product	2.140	0.812	2.069	0.813	2.066	0.816
log of number of rivals ($\phi = 0.2$)	2.265	1.360	-	-	-	-
log of number of rivals ($\phi = 0.15$)	-	-	1.699	1.247	-	-
log of number of rivals ($\phi = 0.10$)	-	-	-	-	1.385	1.115
No. Obs:	177		177		177	
Log-Likelihood	-344.215		-344.683		-344.844	

As described previously, the investment policy is estimated nonparametrically using a Tobit Type I in order to control for the lumpy investment behavior observed in the data. Results are presented on Table

VII. There are no significant changes in the majority of the coefficients for every value of ϕ being considered. The rest of the paper will focus the analysis on the case where $\phi = 0.2$.

The market size state, M_t, was assumed independent of the other states. The log-linear specification proposed for this state can be estimated separately using OLS. The results of this regression are presented on Table VIII. The AR(1) specification in log-market size seems to capture market size dynamics rather well. The inclusion of other regressors (e.g., sum of firm ranks and markup-adjusted inclusive values) would not only bring little explanatory power to the market size transition, but also raise overfitting concerns.

TABLE VIII

MARKET SIZE TRANSITION ESTIMATES

Variable	Estimates	Std. Error
Constant	1.629	0.119
Log of market size	0.804	0.016
R^2	0.995	
No. Obs: 15		

I present the results for the average quality transition function in Table IX. Like for the investment case, the coefficients are hard to interpret. The proposed specification seems to be capturing average quality dynamics rather well. As expected, average quality dynamics seem to depend primarily from past average quality values. Investment impacts positively the average quality, but this effect is not statistically significant. The reason for this low significance and the negative sign of the squared investment coefficient is that the more the firm invests, the more obsolete becomes its product portfolio.

For both the market size and the average quality regressions, the Bera-Jarque normality test didn't reject the normality hypothesis of the error term. So I assume that the error term of both regressions is normality distributed with zero mean.

With all the first-stage estimates at hand, one can simulate the continuation values necessary to estimate ranking benefits and innovation costs. The MLE estimates of these dynamic parameters an presented on Table X. In line with the evidence from the first-stage regressions, I considered $\phi = 0.2$ for the specification of the number of technologically advanced rivals. I also set the discount factor to $\beta = 0.95$.

TABLE IX

AVERAGE QUALITY TRANSITION ESTIMATES

	Estimates	Std. Error
Constant	-18.503	3.485
Investment rate (I_{ft})	0.198	0.124
Squared inv. rate (I_{ft}^2)	-0.001	0.005
Rival's inv. rate $\left(\sum_{i \neq f} I_{it}\right)$	0.038	0.021
average firm quality $(\ln(\upsilon_{ft}))$	2.807	0.557
$(\ln(\upsilon_{ft}))^2$	0.797	0.108
$(\ln(\upsilon_{ft}))^3$	0.081	0.007
$(\ln(\upsilon_{ft}))^4$	0.001	6.89×10^{-5}
Rival quality $\left(\sum_{i \neq f} \upsilon_{it}\right)$	2.486	0.649
Squared rival quality	-0.174	0.053
Cube of rival quality	0.003	0.001
quality cross product $(\ln(\upsilon_{ft})) \times \left(\sum_{i \neq f} \upsilon_{it}\right)$	-0.122	0.090
Squared quality cross product	-0.0034	0.001
Cube of quality cross product	7.35×10^{-5}	1.39×10^{-5}
Market size (M_t)	0.008	0.002
Squared Market size (M_t^2)	-1.27×10^{-6}	5.09×10^{-7}
Firm rank (κ_{ft})	2.030	3.327
Rivals' aggregate rank $\left(\sum_{i \neq f} \kappa_{it}\right)$	1.560	0.549
rank cross product $(\kappa_{ft}) \times \left(\sum_{i \neq f} \kappa_{it}\right)$	-0.263	0.890
Squared rank cross product	-0.251	0.238
Cube of rank cross product	-0.908	0.215
rank and rival quality product $(\kappa_{ft}) \times \left(\sum_{i \neq f} \upsilon_{it}\right)$	0.615	0.289
Squared rank and rival quality product	-0.018	0.007
No. obs	171	
Adj. R^2	0.925	

Estimates suggest that the benefit of being at a certain technological state comands a significant portion of the firm's payoff. For example, the parameters of the benefit function imply that that a technological leader whose rivals have less that 80% of its technological frontier collects an anual benefit of about $897M in 1998 dollars. Estimates also indicate that fixed innovation costs are the most significant portion of the overall innovation cost. However, these costs depend considerably on the firm's frontier positioning. Even though more advanced firms tend to have significant reductions in innovation costs, the industry leader has an additional cost on expanding the industry frontier. Even though this latter effect is not statistically significant, the value obtained for this parameter should be considered in the analysis. The high standard deviation should be primarily due to the fact that only 16 observations (i.e., the total number of periods in the sample) are providing information on this parameter.

TABLE X

DYNAMIC PARAMETER ESTIMATES

	Variable	Estimate	Std. Error
Benefit outside profit terms			
	Firmrank/close rivals ($\kappa_{ft}/N_{\phi=0.2}$)	837.521	116.782
	Squared firmrank/close rivals ($\kappa_{ft}^2/N_{\phi=0.2}$)	59.873	24.934
Innovation cost			
	Constant	128.181	71.651
	Firm rank (κ_{ft})	-81.791	39.915
	Tech. gap on advanced rivals	-14.829	9.934
	Leadership dummy ($1\{\kappa_{ft}=1\}$)	20.912	17.127
	Innovation rate rate (I_{ft})	3.129	1.439
	Square Innovation rate (I_{ft}^2)	1.324	0.789
Standard deviations			
	σ_1	50.612	19.042
	σ_2	0.981	0.263

Another important force which contributes to reductions in fixed costs is the aggregate technological gap between the firm and its most advanced rivals, i.e., $\sum_{i=1}^{F} \mathbf{1}\{\kappa_{it} \geq \kappa_{ft} + \phi\}(\kappa_{it} - \kappa_{ft})$. This is a measure for the ease of imitation for laggards. The intuition is that laggards should be able to lean how to expand their low frontiers by observing the attributes of the supercomputers of their most advanced rivals. Even though it is not as significant as the effects above described, it plays an important role on the ease of frontier expansion.

Finally, the estimates for the variable components of the innovation cost suggest the existence of convexities in frontier expansion.

6 Policy experiments

The motivation for contructing a dynamic model for the supercomputer industry was its ability to simulate counterfactual policy experiments upon recovering the underlying primitives from the data. Several important questions can be answered with the model. For example, what would be the impact of a permanent demand shock on innovation rates? Would a merger between two firms increase equilibrium innovation rates? In this paper, I use the structural model to answer a long-standing question in the Industrial Organization literature: how does technological progress depend on competition? The strategy to obtain the answer consists in two steps. First, I solve the MPNE for investment rates for a given number of competitors for each possible combination of states. Second, taking as given an initial profile of firm computing speed frontiers, I simulate the maximal computing speed available in the industry using the equilibrium investment rates under the different industry sizes. I also compute consumer wellfare and firm profits, and assess its temporal evolution.

Ideally, one would like to perform the proposed experiment with an arbitrary number of firms in order to completely describe the effects of competition in innovation. However, I must restrict attention to a small set of firms and propose a fairly coarse discretization of the state space to make the experiments computationally feasible. For this reason, I conduct the experiment under three different scenarios: monopoly, duopoly and three-firm market. As an initial setup, I consider the three most advanced supercomputer producers in 1997, whose frontiers are presented on Table XI

TABLE XI

FIRM FRONTIERS IN 1997

Firm	Θ_{i1997} (frontier in GFLOPS)
Intel	1338
Cray/SGI	815
Hitachi	368.2

For the state space discretization, I considered four possible values for both demand and average quality. The chosen discretization was

$M = [800, 1600, 2400, 3200]$

$\nu_f = [0.33, 0.66, 1, 1.33]$

For the case of monopoly, there is no need to propose a discretization for the technological state, as it is trivially set to unity. Hence, the total number of states in the state space is 16 (i.e., number of demand states times the number of average quality states). However, a discretization for the tecnological state is necessary for markets with more than one firm. As for other states, I considered four possible values, and the proposed specification was

$\kappa_f = [0.25, 0.50, 0.75, 1]$

Hence, if we let $N \geq 2$ be the number of active firms in the market, then the total number of states in the state space is $4 \times (4 \times 4)^N$. For the case where there are three firms in the market, this corresponds to a total of 16384 states, which imposes a considerable computational burden. However, the fact that each firm's problem satisfies the symmetry and anonymity assumptions allows us to considerably reduce the state space. Indeed, the reduced state space's dimension under such assumptions is $4 \times 16 \times \binom{16+2-2}{2-1} = 8704$. In addition, the symmetry property allows us to restrict attention to the best responses of a single firm. For details on the computation of the MPNE of the game, see Appendix B.

For each initial situation in 1997, I compute profits and consumer welfare by considering the closest neighbor in the discretized space to the true values of M, ν_f and κ_f in that year. Under the logit assumption, consumer welfare is given by

$$CW_t = \frac{M_t}{\alpha} \ln \left(\sum_{f=1}^{F} \exp(-\alpha M k p_f(\kappa_t, \nu_t)) \exp(\delta \kappa_{ft}) \nu_{ft} \right)$$

while total welfare considers not only the sum of profits with CW_t, but also the sum of the differences between benefits from technological state and innovation costs across firms. That is,

$$TW_t = CW_t + \sum_{f=1}^{F} (\Pi_{ft} + \Upsilon_{ft} - C_{ft})$$

The details of the simulation can also be found in Appendix B. For all years except 1997, the values for all entries in Table XII regard averaged values across 20000 simulations. Results suggest that competition significantly encourages technological progress. Even though a similar conclusion can be derived for both consumer and total welfare, it must be mentioned that the welfare gains are decreasing in the number of copetitors. That is, the increases in welfare on moving from monopoly to duopoly are considerably higher compared to the case of moving from duopoly to a three-firm market.

TABLE XII

SIMULATION RESULTS

	1997	1998	1999	2000	2001	2002
Monopoly						
Profit	3315.099	3489.127	3531.561	3614.571	3873.182	4189.067
Consumer welfare	-4149.163	-3783.093	-3567.013	-3278.245	-2997.231	-2818.423
Total welfare	-167.345	-58.82	223.451	612.911	819.423	917.789
Maximal computing speed	1338	1492.061	1789.230	2016.829	2312.678	2567.891
Duopoly						
Firm 1's frontier	1338	1567.231	2401.345	4102.089	7891.341	10613.421
Firm 1 profit	2426.777	2267.89	1918.913	2111.892	2012.123	2178.121
Firm 2's frontier	815	1478.899	2645.212	3912.232	8192.423	9004.561
Firm 2 profit	797.652	1078.912	1511.231	1412.091	1789.991	2009.912
Consumer welfare	1528.039	3489.129	5123.497	7009.324	8213.412	9671.112
Total welfare	5698.558	8910.021	11034.566	13519.112	15013.543	16981.198
Maximal computing speed	1338	1812.411	2916.789	5012.833	9459.341	15329.436
Three-firm market						
Firm 1's frontier	1338	1679.425	3712.321	7092.341	12456.121	18019.642
Firm 1 profit	2006.513	1891.274	1718.291	1801.239	1799.382	1935.924
Firm 2's frontier	815	1618.712	2899.142	6019.348	13513.512	15023.629
Firm 2 profit	699.579	791.231	819.312	891.512	913.123	1016.512
Firm 3's frontier	368.2	901.312	1413.149	5001.432	11013.561	13082.312
Firm 3 profit	324.119	528.519	489.412	621.123	802.231	917.788
Consumer welfare	3620.737	5016.534	6412.521	8812.453	9211.891	10913.451
Total welfare	8221.389	10091.1	12314.567	15012.145	17712.009	18542.612
Maximal computing speed	1338	2221.234	4098.091	10091.521	16716.012	23712.911

7 Conclusions and extentions

This paper proposes a structural model of competition where firms invest in innovation and set constant absolute markups for their products. Strong assumptions imposed in other innovation models proposed in the literature are avoided by building on recent refinements to the dynamic oligopoly games literature. The additional difficulty of modelling innovation in a differentiated products industry is circumvented by using the properties of Nash-Bertrand equilibrium when demand is multinomial logit. The model can be brought to data by using a variant of the two-step estimation method introduced in the literature by Bajari, Benkard and Levin (2006). Identification of the model parameters is based on the product market equilibrium assumptions, orthogonality of selected instruments to unobserved product quality, and optimality conditions on innovation investment. Estimates suggest that firms derive considerable benefits (other than product market profits) from improved technological frontiers. The fact that the incremental payoff from innovating is higher when a firm is in "neck-and-neck" competition with technologically similar rivals (i.e., the "selection effect" introduced in the literature by Aghion, Harris, Howitt and Vickers (2001)) is the source of identification of these benefits.

The paper also quantifies positive externalities of technological states to innovation costs, which indicate that advances on both the firm's maximal computing speed and the ones of its most technologically advanced rivals considerably reduce the fixed costs of innovation investment. The model estimates are also used to assess how does technological progress depend on competition. In line with Aghion, Harris, Howitt and Vickers (2001), I find that increased levels of competition have positive effects on the evolution of supercomputer technology. The fact that firms innovate primarily to escape competition with "neck-and-neck" rivals also impacts positively consumer welfare, even though the magnitude of the improve seems to diminish as extra firms join the market.

Other important questions could be addressed with the model and its estimates. Alternative applications of the results may include, for example, the impact of mergers on the industry's technological frontier, the effects of an exogenous shock in demand on innovation patterns, or the consequences of a given subsidy scheme for innovation behavior. In these cases, the simulation of counterfactuals can be done by making appropriate changes in the state vector. The answer to these and other questions using the proposed model and its estimates is left for future research.

One possible extension of the paper would be to assess what would be the ideal level of competition for this industry be allowing for more than three active firms in the supercomputer market. Even though this analysis is computationally prohibitive under the discretization approach followed in the paper, two alternative methods hold promise on making that extension feasible. The first consists on the function approximation methods, where player's values and policies are represented as infinite-dimensional functions rather than as finite-dimensional vectors. In case where discrete policies such as entry and exit are absent, Doraszelski and Pakes (2005) document a significant reduction in the computational burden. The second alternative method consists on the oblivious-equilibrium approach recently pioneered by Weintraub, Benkard and Van Roy (2005). Under mild assumptions, the oblivious equilibrium closely approximates Markov-Perfect equilibria while allowing for enough computational flexibility to accomodate large industry sizes. The study of these alternatives for answering the question addressed in this paper is left for future research.

Appendix A: Data

For the empirical part of the paper, I collect data on the supercomputer industry from 1990 to 2005 using several sources. The first consists on the TOP500 organization, which collects data on supercomputer worldwide installations twice a year on June and November. The data from this source is publicly available at www.top500.org. The available database contains buyer-specific information about the supercomputers being purchased, except for confidential purchases from selected defense and intelligence agencies. This censoring of data is not likely to be significant, since these purchases seem to represent a low percentage of the overall quantity[20].

One limitation of the TOP500 data is that information is restricted to the 500 installed machines with the highest computing speed according to the Linpack Benchmark at the time of each survey. However, this potential source of selection bias can be neglected for two reasons. The first consists on the definition of supercomputer itself. A computer can only be classified as an HPC if its computing speed is close to the best available one. This time-dependent classification allows me to consider the TOP500 surveys representative of the whole supercomputer market. The second reason consists on the obsolescence of a supercomputer after five years from its introduction. Advances in computing speed become quickly incorporated in other types of computer (mainframes, workstations and eventually laptop and desktop computers), making a supercomputer unlikely to be sold few years after its introduction.

Despite the fact that the TOP500 surveys started on June 1993, I am able to recover information on supercomputer installations since 1990. At the time of that survey, the supercomputer market was yet to experience the remarkable growth in installations during 1990s. Therefore, the earliest available survey of the TOP500 database contains installations dated as far back as 1984. However, there is only a reasonable number of reported installations from 1990 on[21]. The TOP500 dataset includes the identity of the buyer, the type of buyer (Academic, Industry, Government, Classified and Vendor), the name of the supercomputer model, the country and continent where the HPC is installed, the year of installation, quantities purchased per buyer[22], the area of application for the supercomputer, and the specific attributes of the machine. Among those characteristics, one can find two observed quality measures. The first is Rmax, which corresponds to the computing speed of the supercomputer according to the Linpack Benchmark measured in Giga FLOPS. The second measure corresponds to the maximum possible GFLOPS that the HPC can ever process, which is denoted Rpeak. These two computing speed measures are strongly correlated. I found a correlation coefficient of about 0.99 between Rmax and Rpeak. I chose Rmax as the observed quality measure for

[20] The top500 data contains many purchases of classified buyers labelled as "Government Classified", "DoD Classified" and "Defense Classified" along with reported purchases from NSA and FBI. The few available information on confidential purchases indicates that these represent a small percentage of overall supercomputer acquisitions.

[21] Before 1990, the maximal number of installations per year was 21, which is clearly insufficient for obtaining precise estimates.

[22] This information can be inferred directly from the data, since information is provided on an individual-specific basis)

supercomputers for the empirical work.

The other specific details of each supercomputer in the data are system family, number of processors, processor speed, processor brand, interconnect brand, architecture and operating system. Processor speed is measured in Mega-Heartz (MHz), and it determines how many instructions per second can the processor execute. Despite the fact that the top500 surveys take place twice a year on June and November, I decided to consider an yearly frequency for the data. I took this procedure since most supercomputers being installed on year t only appeared on the surveys of year $t+1$ and above. Given that the sample only provides information about the installation year, I pooled the information on all the supercomputers installed on a given year from all surveys. This guarantees a truthful description of the moment of installation of the supercomputer[23].

Unfortunately, the TOP500 dataset provides no reference on the price paid for each supercomputer. Nonetheless, it was possible to collect this information from several different sources. The Transaction Processing Performance Council (TPC), a non-profit organization, publishes detailed information submitted by vendors about prices, performance rating and characteristics of several high-end computers at their website www.tpc.org. Several prices of supercomputers introduced since 1994 were recovered from the TPC Benchmark series TPC-A, TPC-B, TPC-C, TPC-H, TPC-R and TPC-W. Another source where supercomputer prices were found was John McCallum's series on CPU performance. McCallum (2002) uses this series to evaluate price and performance information on several computers from 1944 to 2003[24].

Yearly list price data for several supercomputers was also obtained from the HPC database compilled by the International Data Corporation (IDC)[25]. Even though all these sources only partially covered the price information for the TOP500 supercomputers, it was possible to recover the missing information from searching for the supercomputer models at several other sources. The most representative sources were articles published at Government Computer News (GCN) archives (www.gcn.com), press releases from both manufacturers and buyers, technical reports from NASA[26], Roy Longbottom's Computer Claims data from 1980 to 1996[27], the Federal Procurement Data System, the Government Business Opportunities and the Commerce Business Daily websites[28]. All prices were converted into 1998 constant dollars by using the CPI series from the Bureau of Labor Statistics.

[23] Moreover, the top500 surveys have started in June 1993. It contains information on supercomputers installed in previous years, but it provides no information on the semester of installation. Again, there is only reference to the installation year.

[24] This series is available for download at www.jcmit.com, and the details of its construction can be found in McCallum (2002).

[25] The full database is freely available for download at http://www.hpcuserforum.com/benchmark/

[26] A representative report containing list price informations is the technical report by Saini and Bailey (1996), available for download at http://www.nas.nasa.gov/News/Techreports/1996/PDF/nas-96-018.pdf

[27] Available at http://homepage.virgin.net/roy.longbottom/mips.htm#anchorStart

[28] FPDS procurement data is publicly available for consultancy upon account creation at www.fpdc.gov. Data from the Government Opportunities and from Commerce Business Daily can be obtained by searching on archived awards in their websites. I would like to thank the General Services Administration (GSA) staff of the FOIA and Federal Business Opportunities offices for their helpful explanations on data extraction from these sources.

Market size at a given year is defined as being the number of consumers potentially interested on purchasing a supercomputer. Since supercomputers are durable goods, the intertemporal substitution effects on demand for durables highlighted by Melnikov (2001) must not be neglected on defining the number of potential consumers. As the discrete choice framework proposed for demand does not account for these effects, I will follow the approach of Benkard (2004) by assuming that each buyer in this market optimally reallocates her supercomputer stocks every period, considering both used and new high-performance computers in her choice set. Just like in Benkard's analysis of the market for wide-embodied commercial aircraft, this assumption is acceptable for the case of the supercomputer market. This is because it corresponds to treat supercomputer purchases as rentals, and a considerable number of high-performance computers are installed under rental contracts. Implicit supercomputer rental prices and list prices can be considered proportional, since the maintenance costs for hardware and software for at least three years are included in the latter. Consequently, I defined M_t as the number of of new and used supercomputers in use in year t, which is assumed to be the sum of all supercomputers sold between year $t-5$ and t[29].

Supercomputers which have equal characteristics and are produced by the same firm but sold under a different name were considered to be a single model. Moreover, I assume that two observations in adjacent years represent the same model if (a) they have the same name; (b) their Rmax score, number of processors and processor speed does not change by more than ten percent; and (c) the remaining characteristics (e.g., operative system, interconnectt bandwidth type) are the same. In addition, I excluded both custom supercomputers and vendor systems (i.e., supercomputers that manufacturers produce for their own usage), although these were considered in the computation of technological frontiers and investment rates. Custom systems are not in the same choice set as off-the-shelf models, implying that the logit demand model is only reasonable to describe off-the-shelf purchases. Moreover, the benefits-outside-profits function imposed in the model controls for custom revenues, since the production of custom systems usually implies an expansion of the firm's technological frontier. These procedures yield an unbalanced panel of 2204 model/year observations. Yearly market shares of each supercomputer model are computed by dividing the quantity sold by the market size

[29]This choice was motivated by the fact that the useful lifetime of a supercomputer is five years.

Appendix B: Derivations, proofs and computational details

B1. Derivation of the investment's optimality condition

The firm's problem is

$$V_f(\mathbf{s},\varepsilon_f) = \int\limits_{\varepsilon_{-f}} \underset{\sigma_f(\mathbf{s},\varepsilon_f)}{Max} \left\{ \pi_f(\sigma_f(\mathbf{s},\varepsilon_f), Mkp_{-f}(\mathbf{s}), \mathbf{s}, \varepsilon_f) + \beta E_{\mathbf{s},,\sigma_f,\sigma_{-f}} \left[V_f(\mathbf{s}',\varepsilon_f')|\mathbf{s},\sigma_f(\mathbf{s},\varepsilon_f),\sigma_{-f}(\mathbf{s},\varepsilon_{-f}) \right] \right\} dG(\varepsilon_{-f})$$

where $\mathbf{s} \equiv (M,\hat{s})$, $\hat{s} \equiv (\nu_i,\kappa_i)_{i=1}^F$, $\sigma_f(\mathbf{s},\varepsilon_f) \equiv (Mkp_f(\hat{s}), I_f(\mathbf{s},\varepsilon_f))$, and

$$\pi_f(\sigma_f(\mathbf{s},\varepsilon_f), Mkp_{-f}(\hat{s}), \mathbf{s}, \varepsilon_f) = \Pi_f(Mkp(\hat{s}),\mathbf{s}) + \Upsilon_f(\mathbf{s}) - C(I_f(\mathbf{s},\varepsilon_f),\mathbf{s},\varepsilon_f)$$

where $\Pi_f(Mkp(\hat{s}),\mathbf{s}) = M_t \dfrac{\exp(-\alpha Mkp_f(\hat{s}) + \ln(Mkp_f(\hat{s}))\exp(\delta\kappa_f)\nu_f}{1+\sum\limits_{i=1}^{F}\exp(-\alpha Mkp_i(\hat{s}))\exp(\delta\kappa_i)\nu_i}$

The equilibrium markups were derived from product market competition and are implicitly defined by the following system of F equations

$$Mkp_f(\hat{s}) = \frac{1}{\alpha}\frac{1}{1-q_f(Mkp_f(\hat{s}), Mkp_{-f}(\hat{s}),\hat{s})}, \quad \forall f = 1,...,F$$

where $q_f(Mkp(\hat{s}),\hat{s})$ is the aggregate market share of firm f, which is given by

$$q_f(Mkp(\hat{s}),\hat{s}) = \frac{\exp(-\alpha Mkp_f(\nu,\kappa))\exp(\delta\kappa_f)\nu_f}{1+\sum\limits_{i=1}^{F}\exp(-\alpha Mkp_i(\nu,\kappa))\exp(\delta\kappa_i)\nu_i}$$

We are left to take the first-order condition with respect to investment. Note that, conditional on observed actions, the transition for each firm's technological state, κ_f, is deterministic. The demand state, M, is stochastic but exogenous, while the average quality process is also stochastic but dependent on investment choices. Therefore, the first-order condition for optimality in investment is

$$I_f \left[-\frac{\partial C(I_f,\mathbf{s},\varepsilon_f)}{\partial I_f} + \beta \int\limits_{\varepsilon_{-f}} E_{\mathbf{s},,\sigma_f,\sigma_{-f}} \left(\frac{\partial V_f(\mathbf{s}',\varepsilon_f')}{\partial \kappa_f'}\frac{\partial \kappa_f'}{\partial I_f} + \sum_{i\neq f}\frac{\partial V_f(\mathbf{s}',\varepsilon_f')}{\partial \kappa_i'}\frac{\partial \kappa_i'}{\partial I_f} \right) dG(\varepsilon_{-f}) \right.$$

$$\left. +\beta \int\limits_{\varepsilon_{-f}}\int\limits_{\nu'}\int\limits_{M'}\int\limits_{\varepsilon_f'} V_f(\mathbf{s}',\varepsilon_f') f_{\varepsilon_f'}(\varepsilon_f') f_{M'}(M'|M) \frac{\partial f_{\nu'}(\nu'|\mathbf{s},I_f,\sigma_{-f}(\mathbf{s},\varepsilon_{-f}))}{\partial I_f} d\varepsilon_f' dM' d\nu' dG(\varepsilon_{-f}) \right] = 0$$

Note that the last parcel of this first-order condition can be simplified to

$$\int_{\varepsilon_{-f}} E_{\mathbf{s},,\sigma_f,\sigma_{-f}} \left(V_f(\mathbf{s}', \varepsilon'_f) \frac{\partial \ln \left(f_{\nu'}(\nu' | \mathbf{s}, I_f, \sigma_{-f}(\mathbf{s}, \varepsilon_{-f})) \right)}{\partial I_f} \right) dG(\varepsilon_{-f})$$

after multiplying and dividing by $f_{\nu'}(\nu'|\mathbf{s}, I_f, \sigma_{-f}(\mathbf{s}, \varepsilon_{-f}))$. There is no closed-form solution to the derivatives of the value function with respect to the technological states. However, it is possible to obtain a representation of the first-order condition that is akin to forward simulation. For this purpose, I rewrite the firm problem in sequential problem representation. Using the firm's Bellman equation, this corresponds to

$$V_f(\mathbf{s}_0, \varepsilon_{f0}) = \int_{\varepsilon_{-f0}} \underset{\sigma_f(\mathbf{s},\varepsilon_f)}{Max} E_{\mathbf{s},,\sigma_f,\sigma_{-f}} \left\{ \sum_{t=0}^{\infty} \beta^t \pi_f(\sigma_f(\mathbf{s}_t, \varepsilon_{ft}), Mkp_{-f}(\mathbf{s}_t), \mathbf{s}_t, \varepsilon_{ft}) \right\} dG(\varepsilon_{-f})$$

In what follows, it is convenient to derive the formulas for the derivatives of the future technological states of each firm with respect to current investment decisions. By definition, the transition for κ_f is given by

$$\kappa_{ft+1} \equiv \frac{(1 + I_{ft})\kappa_{ft}}{\max\limits_{i=1,\ldots,F} \{(1 + I_{it})\kappa_{it}\}}$$

By recursive substitution, we have, for $n \geq 1$, the more general formula,

$$\kappa_{ft+n} \equiv \frac{\kappa_{ft} \prod\limits_{j=0}^{n-1} (1 + I_{ft+j})}{\max\limits_{i=1,\ldots,F} \left\{ \kappa_{it} \prod\limits_{j=0}^{n-1} (1 + I_{it+j}) \right\}}$$

so we have

$$\frac{\partial \kappa_{ft+n}}{\partial I_{ft}} = \frac{\kappa_f \prod\limits_{j=1}^{n-1} (1 + I_{ft+j}) \left[\max\limits_{i=1,\ldots,F} \left\{ \kappa_{it} \prod\limits_{j=0}^{n-1} (1 + I_{it+j}) \right\} \right]}{\left[\max\limits_{i=1,\ldots,F} \left\{ \kappa_{it} \prod\limits_{j=0}^{n-1} (1 + I_{it+j}) \right\} \right]^2}$$

$$- \frac{\mathbf{1}\left\{ \kappa_{ft} \prod\limits_{j=0}^{n-1} (1 + I_{ft+j}) = \max\limits_{i=1,\ldots,F} \left\{ \kappa_{it} \prod\limits_{j=0}^{n-1} (1 + I_{it+j}) \right\} \right\} \kappa_{ft}^2 \prod\limits_{j=0}^{n-1} (1 + I_{ft+j}) \prod\limits_{j=1}^{n-1} (1 + I_{ft+j})}{\left[\max\limits_{i=1,\ldots,N} \left\{ \kappa_{it} \prod\limits_{j=0}^{n-1} (1 + I_{it+j}) \right\} \right]^2}$$

and for every $l \neq f$ we have

$$\frac{\partial \kappa_{lt+n}}{\partial I_{ft}} = -\frac{\mathbf{1}\left\{\kappa_{ft}\prod_{j=0}^{n-1}(1+I_{ft+j}) = \max_{i=1,\ldots,F}\left\{\kappa_{it}\prod_{j=0}^{n-1}(1+I_{it+j})\right\}\right\}\kappa_{ft}\kappa_{lt}\prod_{j=1}^{n-1}(1+I_{ft+j})\prod_{j=0}^{n-1}(1+I_{lt+j})}{\left[\max_{i=1,\ldots,F}\left\{\kappa_{it}\prod_{j=0}^{n-1}(1+I_{it+j})\right\}\right]^2}$$

To simplify notation, I define $\Pi_{ft} \equiv \Pi_f(Mkp(\hat{s}_t), \mathbf{s}_t)$, $\Upsilon_{ft} \equiv \Upsilon_f(\mathbf{s}_t)$ and $C_{ft} \equiv C(I_{ft}, \mathbf{s}_t, \varepsilon_{ft})$.Using the sequential representation of the firm's problem, the first order condition on investment is

$$I_f \int_{\varepsilon_{-f}} \left[-\frac{\partial C_{f0}(I_{f0}, \mathbf{s}_0, \varepsilon_{f0})}{\partial I_{f0}}\right.$$

$$+E_{\mathbf{s}_0,,\sigma_f,\sigma_{-f}}\left(\sum_{t=1}^{\infty}\beta^t\left\{\sum_{i=1}^{F}\left(\frac{\partial \Pi_{ft}}{\partial Mkp_{it}}\frac{\partial Mkp_{it}}{\partial \kappa_{it}}\frac{\partial \kappa_{it}}{\partial I_{f0}} + \frac{\partial \Pi_{ft}}{\partial \kappa_{it}}\frac{\partial \kappa_{it}}{\partial I_{f0}} + \frac{\partial \Upsilon_{ft}}{\partial \kappa_{it}}\frac{\partial \kappa_{it}}{\partial I_{f0}} - \frac{\partial C_{ft}}{\partial \kappa_{it}}\frac{\partial \kappa_{it}}{\partial I_{f0}}\right)\right\}\right)$$

$$\left.+E_{\mathbf{s}_0,\sigma_f,\sigma_{-f}}\left(\sum_{t=1}^{\infty}\beta^t\left(\Pi_{ft} + \Upsilon_{ft} - C_{ft}\right)\left\{\sum_{j=1}^{t}\frac{\partial \log\left(f_\nu(\nu_t|\mathbf{s}_{t-1}, \kappa_t, I_{ft}, \sigma_{-f}(\mathbf{s}, \varepsilon_{-f}))\right)}{\partial I_{f0}}\right\}\right)\right] dG(\varepsilon_{-f}) = 0$$

The only terms which depend solely on parameters to be estimated on the second stage are the innovation cost and the benefits outside profits, as well as their derivatives with respect to technological states and investment. The functional forms assumed for these functions are linear in the unknown parameters, so the linearity property also applies to their derivatives.

B2. Value function and derivative dot product decomposition

The fact that the firm's value function and its derivatives are linear allows me to write the expectation terms in the first order condition as a dot product between the vector of second-stage parameters and discounted functions. For the proposed functional forms on innovation costs and technological benefits, we have

$$\frac{\partial C_{ft}}{\partial I_{ft}} = 1\{I_{ft} > 0\}\left(c_4 + 2c_5 I_{ft}\right) + \boldsymbol{\delta}\{I_{ft} > 0\}\left(c_0 + \varepsilon_{1ft} + c_1\kappa_{ft} + \right.$$

$$\left. + c_2\sum_{i=1}^{F}\mathbf{1}\left\{\kappa_{it} \geq \kappa_{ft} + \phi\right\}\left(\kappa_{it} - \kappa_{ft}\right) + c_3\mathbf{1}\left\{\kappa_{ft} = 1\right\} + (c_4 + \varepsilon_{2ft})I_{ft} + c_5 I_{ft}^2\right)$$

$$\frac{\partial C_{ft}}{\partial \kappa_{ft}} = 1\{I_{ft} > 0\}\left(c_1 - c_2\sum_{i=1}^{F}\left(\boldsymbol{\delta}\left(\kappa_{it} - \kappa_{ft} - \phi\right)\left(\kappa_{it} - \kappa_{ft}\right) + \mathbf{1}\left\{\kappa_{it} \geq \kappa_{ft} + \phi\right\}\right)\right)$$

$$\frac{\partial \Upsilon_{ft}}{\partial \kappa_{ft}} = \frac{\rho_1 + 2\rho_2 \kappa_{ft}}{\sum_{i=1}^{F} \mathbf{1}\left\{\kappa_{ft} - \phi \leq \kappa_{it} \leq \kappa_{ft} + \phi\right\}}$$

$$- \frac{\left(\rho_1 \kappa_{ft} + \rho_2 \kappa_{ft}^2\right)\left[\sum_{i=1}^{F} \left(-\boldsymbol{\delta}\left(\kappa_{it} - \kappa_{ft} + \phi\right) \mathbf{1}\left\{\kappa_{ft} + \phi > \kappa_{it}\right\} + \boldsymbol{\delta}\left(\kappa_{ft} - \kappa_{it} + \phi\right) \mathbf{1}\left\{\kappa_{it} + \phi > \kappa_{ft}\right\}\right)\right]}{\left[\sum_{i=1}^{F} \mathbf{1}\left\{\kappa_{ft} - \phi \leq \kappa_{it} \leq \kappa_{ft} + \phi\right\}\right]^2}$$

where $\boldsymbol{\delta}(.)$ denotes Dirac's Delta function. Moreover, for every $i \neq f$, we have

$$\frac{\partial C_{ft}}{\partial \kappa_{it}} = 1\{I_{ft} > 0\}\left(c_2 \boldsymbol{\delta}\left(\kappa_{it} - \kappa_{ft} - \phi\right)\left(\kappa_{it} - \kappa_{ft}\right) + c_2 \mathbf{1}\left\{\kappa_{it} \geq \kappa_{ft} + \phi\right\}\right)$$

$$\frac{\partial \Upsilon_{ft}}{\partial \kappa_{it}} = - \frac{\left(\rho_1 \kappa_{ft} + \rho_2 \kappa_{ft}^2\right)\left[\sum_{i=1}^{F} \left(\boldsymbol{\delta}\left(\kappa_{it} - \kappa_{ft} + \phi\right) \mathbf{1}\left\{\kappa_{ft} + \phi > \kappa_{it}\right\} - \boldsymbol{\delta}\left(\kappa_{ft} - \kappa_{it} + \phi\right) \mathbf{1}\left\{\kappa_{it} + \phi > \kappa_{ft}\right\}\right)\right]}{\left[\sum_{i=1}^{F} \mathbf{1}\left\{\kappa_{ft} - \phi \leq \kappa_{it} \leq \kappa_{ft} + \phi\right\}\right]^2}$$

With this derivations at hand, it is possible to separate the dynamic parameters from discounted functions. For this purpose, let $\boldsymbol{\theta} \equiv (\rho_1, \rho_2, c_0, c_1, c_2, c_3, c_4, c_5)$, and denote $\tilde{\boldsymbol{\theta}} = [1, \boldsymbol{\theta}]$. In addition, define

$$r(I_{f0}, \mathbf{s}_t)\tilde{\theta} = \int_{\varepsilon_{-f}} E_{\mathbf{s}_0, \sigma_f, \sigma_{-f}} \left\{ \sum_{i=1}^{F} \left(\frac{\partial \Pi_{ft}}{\partial Mkp_{it}} \frac{\partial Mkp_{it}}{\partial \kappa_{it}} \frac{\partial \kappa_{it}}{\partial I_{f0}} + \frac{\partial \Pi_{ft}}{\partial \kappa_{it}} \frac{\partial \kappa_{it}}{\partial I_{f0}} + \frac{\partial \Upsilon_{ft}}{\partial \kappa_{it}} \frac{\partial \kappa_{it}}{\partial I_{f0}} - \frac{\partial C_{ft}}{\partial \kappa_{it}} \frac{\partial \kappa_{it}}{\partial I_{f0}} \right) \right\} dG(\varepsilon_{-f})$$

$$w(I_{f0}, \mathbf{s}_t)\tilde{\theta} = \int_{\varepsilon_{-f}} E_{\mathbf{s}_0, \sigma_f, \sigma_{-f}} \left\{ \left(\Pi_{ft} + \Upsilon_{ft} - C_{ft}\right) \left\{ \sum_{j=1}^{t} \frac{\partial \log\left(f_\nu(\nu_t | \mathbf{s}_{t-1}, \kappa_t, I_{ft}, \sigma_{-f}(\mathbf{s}, \varepsilon_{-f}))\right)}{\partial I_{f0}} \right\} \right\} dG(\varepsilon_{-f})$$

Recall that the first-order condition on investment can be written as

$$I_f \left[-\frac{\partial C(I_f, \mathbf{s}, \varepsilon_f)}{\partial I_f} + \beta(\hat{R}(\mathbf{s}, I_f) + \hat{W}(\mathbf{s}, I_f))\tilde{\theta} \right] = 0$$

The dot product representation follows by noting that $\hat{R}(\mathbf{s}, I_f) = \sum_{t=1}^{\infty} \beta^t r(I_f, \mathbf{s}_t)$ and $\hat{W}(\mathbf{s}, I_f) = \sum_{t=1}^{\infty} \beta^t w(I_f, \mathbf{s}_t)$. A similar decomposition applies to the value function itself, yielding the representations $\hat{Y}_\sigma(\mathbf{s})$ and $\hat{Y}_{\sigma_{-f} I=0}(\mathbf{s})$) for MPNE and zero-investment, respectively.

B3. Forward simulation details

Given any initial state $\mathbf{s}_0$ and investment choice I_{f0} each component of the vector $\hat{R}(\mathbf{s}_0, I_{f0})$ (and vector $\hat{W}(\mathbf{s}_0, I_{f0})$) for firm f can be simulated by repeating K times the following steps:

- For each of firm f rivals, draw a random variable from the Normal distribution with zero mean and variance equal to the one implied by the estimated Tobit in the first-stage policy estimation. Compute the rival's investment choices conditional on the draw and observed states.
- Compute $r(I_{f0}, \mathbf{s}_0)$ and $w(I_{f0}, \mathbf{s}_0)$
- Compute each firm's new frontier rank and draw both a new demand state and average qualities for the active firms;
- Restart from the first step up to T^* times (i.e., until you reach the period at which discounted payoffs are assumed not significant).
- Discount the results with $\beta^\tau, \tau = 1, .., T^*$ and sum across τ.
- Compute the average across the K simulations;

As this procedure can be carried out for an arbitrary initial state, one can easily obtain the simulated terms for every firm and period. For the empirical application, I considered T^*

B4. Counterfactual experiment details

Given the estimated transitions from the first-step estimates, it is possible to assign a probability to each state by integrating its transition density over a given region of integration. I choose the midpoints between states as the integration limits for each state's probability,. In case the state in question is not bounded above (namely for M and ν_f), the upper integration limit was infinity. Conditional on current states and on firm policies, the state transition probabilities are evaluated by integrating each state's transition density on the chosen integration region. Following Ryan (2004), a small degree of uncertainty was imposed for the deterministic state κ_n for computational convenience. In particular, the following transition probability was imposed:

$$\Pr(\kappa'_f | \kappa, I_f, I_{-f}) = \int_{\Lambda_1}^{\Lambda_2} dH(u)$$

where $H \tilde{} U \left[\frac{(1+I_{ft})\kappa_{ft}}{\max\limits_{i=1,\dots,F} \{(1+I_{it})\kappa_{it}\}} - \eta, \frac{(1+I_{ft})\kappa_{ft}}{\max\limits_{i=1,\dots,F} \{(1+I_{it})\kappa_{it}\}} + \eta \right]$ and η is the distance between two adjacent nodes in the state space.

The integration limits correspond to

$$\Lambda_1 = \frac{(1+I_{ft})\kappa_{ft}}{\max\limits_{i=1,\ldots,F}\{(1+I_{it})\kappa_{it}\}} - \eta/2$$

$$\Lambda_2 = \frac{(1+I_{ft})\kappa_{ft}}{\max\limits_{i=1,\ldots,F}\{(1+I_{it})\kappa_{it}\}} + \eta/2.$$

This specification has the advantage improve convergence considerably by guaranteeing smooth best response functions of a firm to its rivals while nesting the deterministic transition in the case of an infinitely fine state space.

The exchangeability and symmetry assumptions allow us to convert rival policies into own firm policies, one can restrict attention to a single firm's problem in the computations. I will closely follow the exposition of Collard-Wexler (2006) to clarify the usefulness of these two assumptions. Suppose there are three firms in the market. Exchangeability means that, for the perspective of a given firm, it is irrelevant the identity of the opponents. For example, suppose there is only one sate per firm, which can assume three values: 0.1, 0.2 and 0.3. Under exchangeability, firm 1 will get exactly the same payoff both when firm 2's state is 0.1 and firm 3's state is 0.2 and vice-versa. This applies for every possible state of firm 1. Hence, I am able to reduce the state space by grouping permutations of the rival's states as being one single state and adding their probabilities. In the previous example, this corresponds to considering the pairs of rival states (0.1,0.2) and (0.2,0.1) as a single state and sum the probabilities of these states, for every state of firm 1. I denote $\mathbf{s}^*$ as being the reduced state space for each firm.

The symmetry assumption implies that the payoff structure of each firm is the same. For example, the payoff firm 1 receives when its state is 0.2 and the opponents' states are 0.1 and 0.3 must be the same as the payoff that firm 3 obtains when it state is 0.2 and its opponents' states are 0.1 and 0.3. Similar implications apply to the policy functions. Therefore, I can restrict attention to a single firm's problem, since the policies of its opponents can be written in terms of the firm's own policies.

In order to compute the MPNE of the game, it is necessary to restrict attention to the ex-ante value function of each firm. This requires integration of all private information of each firm's problem, which yields the following integrated Bellman equation

$$V_f(\mathbf{s}) = \int_{\varepsilon} \underset{\tilde{\sigma}_f(\mathbf{s},\varepsilon_f)}{Max} \left\{ \pi_f(\sigma_f(\mathbf{s},\varepsilon_f),\mathbf{s},\varepsilon_f) + \beta \int V_f(\mathbf{s}')dP(\mathbf{s}'|\mathbf{s},\tilde{\sigma}_f(\mathbf{s},\varepsilon_f),\tilde{\sigma}_{-f}(\mathbf{s},\varepsilon_{-f})) \right\} dF(\varepsilon)$$

It can be shown that this is a contraction by checking the conditions of Blackwell's Theorem. Therefore, the optimal (integrated) investment policies yield an unique value function for each firm. Since the decisions of a firm affect the states of other firms, a policy function iteration algorithm was preferred to value function iteration proposed by Pakes and McGuire (1994). The algorithm can be outlined as follows.

- Start with an initial guess of the firm's (integrated) investment policies. Denote it $I_{f0}(\mathbf{s}^*)$, where $\mathbf{s}^*$ corresponds to a state vector in the reduced state space. Use the same guess for the remaining firm's policies

- Compute the firm's value function, $\Psi_{f0} = \left[I_{S^* \times S^*} - \beta \Pr(\mathbf{s}^{*\prime} | \mathbf{s}^*, I_{f0}(\mathbf{s}^*)\right]^{-1} \Gamma_n(\mathbf{s}^*, I_{f0}(\mathbf{s}^*))$, using the guesses of the previous step. $\Gamma_n(\mathbf{s}^*, I_{f0}(\mathbf{s}^*))$ corresponds to the vector of the firm's flow payoff, while $I_{S^* \times S^*}$ is an identity matrix of dimension $S^* \times S^*$. This is the policy evaluation step.

- Compute $I_{f1}(\mathbf{s}^*) \equiv \underset{I_{f1}(\mathbf{s}^*)}{\arg\max} \{\Gamma_n(\mathbf{s}^*, I_{f1}(\mathbf{s}^*)) + \beta \Pr(\mathbf{s}^{*\prime} | \mathbf{s}^*, I_{f1}(\mathbf{s}^*)) \Psi_{f0}\}$. subject to the (equality) constraints implied by symmetry and exchangeability to the rival's policies. This is the policy improvement step

- If, for some given tolerance limit d, we have $||I_{f1}(\mathbf{s}^*) - I_{f0}(\mathbf{s}^*)|| < d$, stop the algorithm and take $I_{f1}(\mathbf{s}^*)$ as the solution. Otherwise go back to step 1 using $I_{f1}(\mathbf{s}^*)$ instead of $I_{f0}(\mathbf{s}^*)$.

Upon solving for the MPNE under one, two and three-firm markets, it is possible to simulate the evolution of the industry for the variables of interest. The simulation process can be described as follows.

- For each firm f, read the vector of states in 1997[30]

- Compute the MPNE investment rate choices by reading the solution to the DP problem at the state vector;

- Given the MPNE investment choices, compute the equilibrium profits, consumer welfare and total welfare.

- Update the state vector by sampling a draw from the transition probabilities evaluated at the investment choices and current states;

- Compute the new frontier of each player in every slot. For every firm f, this is simply $(1 + I_{f1997}) \times \Theta_{f1997}$. Compute the maximal new frontier in the industry.

- Restart all this procedure, taking the sampled state vector as the 1998 state vector. Repeat all these procedures until 2002.

- Repeat these steps L times;

[30] In practice, the state vector to be considered will be the closest neighbor to the 1997 state vector due to the state space discretization.

- Compute each year's expected industry frontier, firm profits and welfare measures by averaging each maximum frontier across the L simulations for every given year.

For the policy experiments, I chose L =20000 and $\beta = 0.95$.

References

[1] Ackerberg, D., C. L. Benkard, S. Berry and A. Pakes (2006): "Econometric Tools for Analyzing Market Outcomes", working paper prepared for J. Heckman and E. Leamer (eds), *Handbook of Econometrics*, Volume 6.

[2] Aghion, P., C. Harris, P. Howitt and J. Vickers (2001): "Competition, Imitation and Growth with Step-by-Step Innovation", *The Review of Economic Studies*, Vol. 68, No. 3, pp 467-492.

[3] Aguirregabiria, V. and P. Mira (2006): "Identification and Estimation of Dynamic Discrete Games", forthcoming *Econometrica.*

[4] Anderson, S. A. Palma and J. Thisse (1992): "Discrete choice theory of product differentiation", The MIT Press, Cambridge, Mass.

[5] Collard-Wexler, A. (2006): "Plant Turnover and Demand Fluctuations in the Ready-Mix Concrete Industry", Stern School of Business, New York University, mimeo.

[6] Athey, S. and A. Schmutzler (2001): "Investment and Market Dominance", *The RAND Journal of Economics*, Vol. 32 pp 1-26.

[7] Bailey, D. and Saini, S. (1996): "NAS Parallel Benchmark (Version 1.0) Results 11-96", NASA Ames Research Center, Report NAS-96-18.

[8] Bajari, P., C. L. Benkard and j. Levin (2006): "Estimating Dynamic Models of Imperfect Comeptition", forthcoming *Econometrica.*

[9] Benkard, C.L. (2004): "A Dynamic Analysis of the Market for Wide-bodied commercial aircraft", *The Review of Economic Studies*, Vol. 71 pp 581-611.

[10] Benkard, C.L., B. Van Roy and G. Weintraub (2005): "Markov Perfect Industry Dynamics with Many Firms", Stanford University, mimeo.

[11] Beresteanu, A. and P. Ellickson (2006): "The Dynamics of Retail Oligopoly", Duke University, mimeo.

[12] Berry, S. (1994): "Estimating Discrete-Choice Models of Product Differentiation", *The RAND Journal of Economics*, Vol. 25, No. 2, pp 242-262.

[13] Berry, S., J. Levinsohn and A. Pakes (1995): "Automobile Prices in Market Equilibrium", *Econometrica*, Vol. 63, No. 4, pp 841-890.

[14] Berry, S., M. Ostrovsky and A. Pakes (2005): "Simple Estimators for the Parameters of Dynamic Games (with Entry/Exit Examples)", NBER Working Paper No. 10506.

[15] Bresnahan and Greenstein (1999): "Technological Competition and the Structure of the Computer Industry", *The Journal of Industrial Economics*, Vol. 47, pp 1-40.

[16] Bresnahan, T., S. Stern and M. Trajtenberg (1997): "Market Segmentation and the Sources of Rents from Innovation: Personal Computers in the Late 1980s", *The RAND Journal of Economics*, Vol. 28 pp S17-S44.

[17] Bresnahan, T. and M. Trajtenberg (1995): "General Purpose Technologies 'Engines to Growth'?", *The Journal of Econometrics*, Vol. 65, pp 83-108.

[18] Caplin, A. and B. Nalebuff (1991): "Aggregation and Imperfect Competition: On the Existence of Equilibrium", *Econometrica*, Vol. 59, No. 1, pp 26-61.

[19] Carranza, J. (2006): "Product Innovation in Durable Goods Markets", University of Wisconsin-Madison, mimeo.

[20] Dongarra, J., Luszczek and A. Petitet (2001): "The LINPACK Benchmark: Past, Present and Future", University of Tennessee, mimeo.

[21] Dongarra, J. (2006): "Performance of Various Computers Using Standard Linear Equations Software", University of Tennessee, technical report.

[22] Doraszelski, U. and M. Satterthwaite (2005): Foundations of Markov-Perfect Industry Dynamics: Existence, Purification and Multiplicity", Northwestern University, technical report.

[23] Doraszelski, U. and A. Pakes (2006): "A Framework for Applied Dynamic Analysis in IO", Harvard University, mimeo.

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

[25] Evans, D. and R. Schmalensee (2002): "Some Economic Aspects of Antitrust Analysis in Dynamically Competitive Industries", in *Innovation Policy and the Economy*, by A. Jaffe, J. Lerner and S. Stern (eds), MIT Press.

[26] Gowrisankaran, G. (1999): "A Dynamic Model of Endogenous Horizontal Mergers", *The RAND Journal of Economics*, Vol. 30 pp 56-83.

[27] Graham, S., M. Snir and C. Patterson (2005): "Getting up to speed: the future of Supercomputing", the National Academies Press.

[28] Green, J. and S. Scotchmer (1995): "On the Division of Profit in Sequential Innovation", *The RAND Journal of Economics*, Vol. 32 pp 20-33.

[29] Hendel, I. (1999): "Estimating Multiple-Discrete Choice Models: An Application to Computerization Returns", *The Review of Economic Studies*, Vol. 66 pp 423-446.

[30] Hendel, I. and A. Nevo (2006): "Measuring the Implications of Sales and Consumer Inventory Behavior", forthcoming *Econometrica.*

[31] Hotz, J., R. Miller, S. Sanders and J. Smith (1994): "A Simulator Estimator for Dynamic Models of Discrete Choice", *The Review of Economic Studies*, Vol. 61, No. 2, pp 265-289.

[32] Khanna, T. (1995): "Racing Behavior Technological Evolution in the High-end Computer Industry", *Research Policy*, Vol. 24, pp 933-958.

[33] Lucas, R. and N. Stokey (1989): "Recursive Methods in Economic Dynamics", Harvard Univeristy Press.

[34] Mazzeo, M. (2002): "Product Choice and Oligopoly Market Structure", *The RAND Journal of Economics*, Vol. 33, No. 2, pp 1-22.

[35] McCallum, J. (2002): "Price-Performance of Computer Technology", in *The Computer Engineering Handbook*, by V. Oklobdzija (ed), CRC Press.

[36] McFadden, D. (1981): "Econometric Models of Probabilistic Choice", in *Structural Analysis of Discrete Choice with Econometric Applications*, by D. McFadden and C. Manski (eds), MIT Press.

[37] Melnikov, O. (2001): "Demand for Differentiated Durable Products: The Case of the U.S. Computer Printer Market", Cornell University, mimeo.

[38] Nevo, A. (2000): "Mergers with Differentiated Products: The Case of the Ready-to-Eat Cereal Industry", *The RAND Journal of Economics*, Vol. 31 pp 395-421.

[39] Nevo, A. (2001): "Measuring Market Power in the Ready-to-Eat Cereal Industry", *Econometrica*, Vol. 69, No. 2, pp 307-342.

[40] Nevo, A. and F. Rossi (2006): "An approach for extending dynamic models to settings with multi-product firms", Northwestern University, mimeo.

[41] Pakes, A. and P. McGuire (1994): "Computing Markov-Perfect Nash Equilibria: Numerical Implications of a Dynamic Differentiated Product Model", *The RAND Journal of Economics*, Vol. 25 pp 555-589.

[42] Rust, J. (1987): "Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher", *Econometrica*, Vol. 55, No. 5, pp 999-1033.

[43] Ryan, S. (2006): "The Costs of Environmental Regulation in a Concentrated Industry", MIT, mimeo.

[44] Rossi, F. (2006): "An Empirical Model of Dynamic Attribute-Space Competition", Kellogg School of Management, Northwestern University, mimeo.

[45] Scotchmer, S. (1996): "Protecting Early Innovators: Should Second-Generation Products be Patentable?", *The RAND Journal of Economics*, Vol. 27 pp 322-331.

[46] Segal, I. and M. Whinston (2004): "Antitrust in Innovative Industries", Stanford and Northwestern Universities, mimeo.

[47] Seim, K. (2005): "An Empirical Model of Firm Entry with Endogenous Product-Type Choices", forthcoming *The RAND Journal of Economics*.

[48] Song, Minjae (2006): "A Dynamic Analysis of Cooperative Research in the Semiconductor Industry", School of Economics, Georgia Institute of Technology, mimeo.

[49] Trajtenberg, M. (1989): "The Wellfare Analysis of Product Innovations, with an Application to Computed Tomography Scanners", *The Journal of Political Economy*, Vol. 97 pp 444-479.