NBER WORKING PAPER SERIES

MARKET STRUCTURE AND INNOVATION:
A DYNAMIC ANALYSIS OF THE GLOBAL AUTOMOBILE INDUSTRY

Aamir Rafique Hashmi
Johannes Van Biesebroeck

Market Structure and Innovation: A Dynamic Analysis of the Global Automobile Industry
Aamir Rafique Hashmi and Johannes Van Biesebroeck
NBER Working Paper No. 15959
May 2010
JEL No. C73,L13,L62,O31

ABSTRACT

We study the relationship between market structure and innovation in the global automobile industry from 1982 to 2004 using the dynamic industry framework of Ericson and Pakes (1995). Firms optimally choose a continuous level of innovation in a strategic and forward-looking manner, while anticipating the possibility of future mergers. We show that our estimated model predicts the data well and that changes in the modeling assumptions have a predictable effect on the key dynamic parameter -- the cost of innovation. In terms of the relationship between market structure and innovation, we find that: (1) At the firm level, there is a weakly positive relationship between a firm's price-cost margin and its innovation intensity; (2) There is no relationship between competition and innovation at the industry level in the steady state. As the industry goes through a consolidation phase, the relationship is negative if competition is measured by the inverse of markups and positive if it is measured by the inverse of concentration; (3) A key determinant of a firm's innovation intensity is its relative position in the industry in terms of knowledge stock.

Aamir Rafique Hashmi
Department of Economics
National University of Singapore
1 Arts Link
AS2/05-08
Singapore 117570
aamir@nus.edu.sg

Johannes Van Biesebroeck
Centre for Economic Studies
K.U.Leuven
Naamsestraat 69
3000 Leuven
Belgium
and NBER
jo.vanbiesebroeck@econ.kuleuven.be

1 Introduction

Schumpeter [1942] famously advanced the argument that monopoly is more conducive to innovation than highly competitive markets. An extensive literature investigating the effect of market structure on innovative activity followed, but it has proven difficult to identify robust empirical results. Cohen and Levin [1989] highlight several methodological difficulties that have plagued empirical work. A major difficulty, as a number of theoretical studies have demonstrated, is that the competition-innovation relationship is only monotonic under restrictive circumstances (Gilbert [2006]).

One reason for this is the opposing impact of the 'replacement' and 'efficiency' effects. The former leads to lower innovation incentives for a monopolist that has existing profits at stake. The latter leads to lower innovation incentives in more competitive situations as competition lowers aggregate industry profits. Aghion et al. [2005] demonstrate a nonlinear 'inverted-U' pattern between competition and innovation, confirming the suggestive results of Scherer [1967] and Levin et al. [1985] in an explicit model of firm optimizing behavior. At low levels of competition, more competition is associated with higher innovation, but the efficiency effect starts to dominate at higher levels of competition.

A second reason for the paucity of robust results is the endogeneity of market structure. Vives [2008] shows some robust effects in a set of empirically relevant cases, but results remain sensitive to the assumption of an exogenous market structure versus endogenous entry. A related problem for empirical work is reverse causality from innovation to market structure. The issue has been ignored in much of the literature, although a few studies have exploited quasi-natural experiments that feature exogenous changes in market structure. For example, Carlin et al. [2004] follow firms from transition countries over time as competition is suddenly introduced after abandoning the communist model. Aghion et al. [2005] instrument competition changes with policy variables associated with the integration of the U.K. economy into the European market.

Several other factors can make it impossible to identify a stable empirical relationship between innovation and market structure. The exact way innovation enters the model matters greatly. Opposite effects have been demonstrated for product versus process innovations (Boone [2000a], Vives [2008]), for discrete innovations that make an existing technology obsolete versus those that do not (Gilbert [2006]), and for cases with or without complementarities between different innovation decisions (Kretschmer et al. [2009]).

Identifying the relationship from cross-sectoral variation is complicated by variation in profit opportunities and innovation costs across industries. Market structure variables often turn insignificant if sufficient controls are included (Gilbert [2006]). Blundell et al. [1999] rely on differences

between firms within an industry instead, but heterogeneity between firms is likely to be important as well (Boone [2000b]).

All of the above factors are problematic for reduced form studies that regress a measure of innovation on a measure of competition or market structure. Estimates will be sensitive to functional form assumptions and to the controls included. In contrast, we propose to study the same relationship in an explicit dynamic model of strategic decision making, an approach urged by Cohen and Levin [1989], Gilbert [2006], and Sutton [2008]. Instead of looking for relationships that are monotone throughout an entire data set, we can focus on local effects and the model will also suggest which controls are crucial. Two-way interactions between market structure on innovation can be studied if the primitives of the model are estimated. In addition, one could even study how the market equilibrium and innovation choices change if a parameter is changed exogenously.

We apply our model to the global automotive industry as it is a highly innovative industry, both in terms of R&D expenditure and patents granted, and it has recently experienced large changes in market structure. The widely varying fortunes of different firms have given rise to large organic changes in market share. In addition, mergers have caused discrete changes in market structure directly affecting more than two thirds of the firms in the sample.

Our principal objective is to study the response of firms in terms of their innovative activity to changes in the level of competition—a combination of market structure and the level of innovation of all market participants. We do this in an explicit model of industry dynamics that incorporates several important features of the automobile industry.

In our model, each firm produces one differentiated product that is characterized by two quality dimensions. The observable dimension of quality is proxied by the patent stock of a firm, which can be augmented through R&D investments. At the same time, R&D expenditures potentially lower the marginal cost of production. The second dimension of quality, capturing something like popularity of a new design, evolves exogenously and is unobservable to the econometrician. The market share of a firm depends on the relative quality of its product and the price, which is set strategically in each period. R&D investment is modeled as a strategic decisions: a firm takes the actions of its rivals and their possible future states into account.

After each period, there is a probability that any two market participants merge. Mergers are introduced as exogenous events, but forward-looking firms take them into account when making optimal innovation decisions. Because knowledge stocks, the cumulative result of past innovation, influence the division of a newly created firm's value over the two merging firms, the mere possibility of a future merger influences innovation incentives for each firm and the interactions between firms.

A few other papers study the interrelation between innovation and market structure in a dy-

namic model of strategic interaction. Goettler and Gordon [2009] study the microprocessor industry and explicitly incorporate the durable nature of the good by making demand and price setting dynamic as well. As the industry has long been a duopoly of Intel and AMD, the impact of market structure on innovation can only be assessed by a counterfactual analysis of monopoly innovation, relying on the primitives estimated from the duopoly model. Their finding that as a monopoly Intel would innovate more depends crucially on the durable nature of the good, upgrades are necessary to stimulate demand. It only holds if consumers value quality highly and are relatively price insensitive.

Xu [2008] analyzes innovation decisions in the Korean electric motor industry. In addition to the cost of R&D, he also estimates R&D spillovers, adjustment costs of physical investment, and the distribution of plant scrap values. As he uses the oblivious equilibrium concept of Weintraub et al. [2008], there is no strategic interaction between plants in the innovation decision. Plants only optimize relative to a stable industry state. Finally, Siebert and Zulehner [2008] study the reverse question of ours, i.e. how innovation affects market structure. In their study of the DRAM industry, they estimate the evolution of sunk entry costs as the innovation intensity and market demand increase over time. Through their effect on entry and exit, these costs determine equilibrium market structure.

Compared to these other studies we introduce two innovations. First, we model a continuous control variable, innovation, in a dynamic context. Firms optimize dynamically, taking actions of competitors into account. We use the two-step estimation strategy of Bajari et al. [2007] for this dynamic game, which does not require to solve for the equilibrium—an impossible task for the 23 active firms in our sample.[1] A second innovation is the introduction of mergers. It greatly increases the computational burden as the linearity of the Bajari et al. [2007] approach is lost, but it does introduce additional and exogenous variation in market structure. Through their impact on market structure, mergers directly affect innovation incentives. Through their impact on the future payoff of current innovations, they have an additional indirect effect.

Another contribution of our study is to provide a parsimonious model of competition in the global automobile industry. We find that the technological knowledge of a firm, even after controlling for other observable product characteristics, is a strong determinant of the demand for vehicles. This finding highlights the importance of innovation in this industry.

---

[1]This estimator has only been used in a few other empirical applications. Examples are Ryan [2006], who studies the effect of environmental regulation on the cement industry, and Ellickson and Misra [2008] who study supermarket pricing. Ching [2008] modified the estimator, simultaneously estimating the demand and policy functions, to study demand dynamics in prescription drugs after patent expiry. Jenkins et al. [2004] also modified the estimator, allowing the dynamic coefficients to enter nonlinearly, to estimate network effects in the market for internet browsers.

Parameter estimates for the primitives of the model are of plausible magnitudes. In the benchmark model, the key dynamic parameter—the cost of innovation—is estimated at $21.8 million per patent application. Changes in the modeling assumptions have a predictable effect on this estimate. For example, alternative demand estimates that lead to a lower price elasticity boost the profitability of patents. To still match the observed patenting rate in the data, the cost of innovation is estimated 30-40% higher.

While the absence of robust monotone effects is a problem for reduced form studies, it is an advantage for structural approaches. The richness of the interrelations between innovation and other primitives, such as demand or cost, and other firm choices, such as price setting or entry, can give rise to a variety of patterns. In terms of the relationship between market structure and innovation, we use forward simulation of the model to illustrate that: (1) At the firm level, there is a weakly positive relationship between a firm's price-cost margin and its innovation intensity; (2) There is no relationship between competition and innovation at the industry level in the steady state. As the industry goes through a consolidation phase, the relationship is negative if competition is measured by the inverse of markups and positive if it is measured by the inverse of concentration; (3) A key determinant of a firm's innovation intensity is its relative position in the industry in terms of knowledge stock.

The rest of the paper is organized as follows. Section 2 provides background information on innovation and the evolving market structure in the global automotive industry. The supply and demand side of the model as well as the Markov Perfect equilibrium concept we rely upon are introduced in Section 3. Section 4 introduces the data. The two-step estimation methodology and the coefficient estimates are discussed in Section 5, where we also examine the robustness of our estimates (in Section 5.3). The estimated model is used to analyze the equilibrium interaction between innovation and market structure in Section 6 and Section 7 concludes.

2   The Automotive Industry

The automotive industry is well-suited to investigate the interaction between innovation and market structure in a strategic context. Demand estimates—see for example Berry et al. [1995] and Goldberg [1995]—reveal that markups over marginal costs tend to be large, consistent with the view that fixed costs are important in this industry. Innovation is an important source of product differentiation as firms' competitive positions are improved through higher product quality, greater reliability and the introduction of new product features. In addition, the industry is the poster child for the importance of process innovations that continually boost productivity (Van

Biesebroeck [2003]).

Developing and producing automobiles is a highly research intensive activity: in 2006, more than 13% of all R&D in the OECD was spent in ISIC industry 34 'Motor Vehicles', more than in any other industry. Statistics in Table 1 highlight the importance of automotive R&D in the five most research intensive economies. Except for the U.S. where it is fourth, the industry is first or second in terms of R&D spending in all countries. The top 13 firms in the auto industry spent more than 55 billion dollars on R&D in 2005. The industry is also a heavyweight on the output side of the innovation process. In the last 25 years, those same 13 firms were awarded more than 50,000 patents by the U.S. patent office.[2]

The industry is concentrated worldwide, making it likely that firms will take actions of competitors into account when deciding on their own innovative activities. In 2004, more than 95% of all vehicles were sold by the 13 largest firms, which were active in all major regions of the world.[3]

The global automobile industry has seen significant consolidation over the last few decades. Many of the industry giants have found it beneficial to join hands with some of their former rivals. The mergers between Daimler-Benz and Chrysler and between Hyundai and Kia, the association between Renault and Nissan and the takeover of Mazda, Jaguar and Volvo by Ford are but a few examples of this consolidation.[4] On the one hand, this consolidation is the result of increased competition and high research intensity, which has made it harder for smaller firms to survive on their own. On the other hand, consolidation has an impact on the intensity of competition as well. With fewer firms around, standard economic models would predict less competition. However, the emerging groups are more evenly matched in terms of research intensity and they are more likely to compete head-on over their entire product range.

At the aggregate level, the equilibrium relationship between market structure and aggregate innovation suggests that more competition is associated with less innovation. The raw data, in the top-left panel of Figure 1, shows a negative relationships between the number of active firms and the total number of patents granted in the industry. Patenting was a lot more intense in the last years of the sample period, with 13 active firms, than in the earlier years, with 23 active firms. The negative relationship is only slightly less pronounced if patents divided by revenue is used as

---

[2]5% of all patents filed in the European patent office (by E.U. applicants) are in the narrow IPC category B60 for "vehicles in general", which is only one of 127 categories and contains only a subset of motor vehicle related innovations; the corresponding fraction at the U.S. patent office is 3%.

[3]We do not distinguish between minority share holdings and outright control. E.g. During our sample period, Mazda is counted as part of the Ford group, even though Ford Motor Co. never held more than 33.4% of Mazda's shares.

[4]The deterioration of the automotive market in North America and Europe has led to the reversal or reconfiguration of many of these mergers, but that happened after the end of our study period.

measure of innovation, as shown in the bottom-left panel.[5]

It is well known however, that the rate of patenting has increased over time in all industries, especially after 1984 (Hall [2004]). If we purge the innovation variable on the vertical axis from a flexible time trend with a cubic series and plot the residual instead, the pattern is not as clear-cut. Patenting now appears to be stronger in early years and the relationship between the number of firms and the total number of patents becomes U-shaped. Using patents by revenue as innovation measure, the relationship remains downward-sloping—less innovation when more firms are active— but now at a decreasing rate.

The aggregates hide important variation at the micro-level. Not all firms patent equally. While some firms choose to patent a little or not at all, others have consistently high rates of innovation and invariably file many patent applications. Moreover, partly as a result of the mergers, there is some convergence over time in the patenting rates across firms. The relationship between the size of each firm, in terms of revenue, and patent stock that each chooses to hold is depicted in Figure 2 for four different years. In each period, the relationship is inverted-U shaped. On average the smallest firms patent the least and hold the smallest patent stocks. The most innovative firms are the ones in the middle of the size distribution.[6]

This pattern would be even more pronounced if patents by revenue would be used on the vertical axis. It is consistent with the evidence for U.K. manufacturing firms in Aghion et al. [2005], but in order to gauge whether the mechanisms leading up to this are also consistent with their interpretation, we need a model of firm behavior and industry competition.

3   The Model

Our model follows Ericson and Pakes [1995]. We keep it as simple as possible to introduce a novel aspect in the dynamic game: the mergers.

There are n firms, each producing a single differentiated vehicle. Firms differ in their technological knowledge, $\omega_j \in \mathbb{R}^+$, and in a firm-specific quality index, $\xi_j \in \mathbb{R}$. For the industry as a whole, the vectors containing ω and ξ are denoted by $\mathbf{s}_\omega$ and $\mathbf{s}_\xi$ and $\mathbf{s} = \{\mathbf{s}_\omega, \mathbf{s}_\xi, m\}$, where $m \in \mathbb{N}$ is the number of consumers in the market. The vectors of ω and ξ excluding firm j are $\mathbf{s}_\omega^{-j}$ and $\mathbf{s}_\xi^{-j}$, and $\mathbf{s}^{-j} = \{\mathbf{s}_\omega^{-j}, \mathbf{s}_\xi^{-j}\}$.

---

[5]Revenue numbers are relative to GM, obtained by multiplying the number of vehicles sold by the price relative to GM (see below).

[6]The dashed line depicts the average statistical relationship between firm size and patent stock estimated over the entire sample of all firm-year observations. The solid lines re-estimate this relationship in each year.

Time is discrete. At the beginning of each period, firms observe **s** and decide on price and investment in new knowledge. While pricing decisions are static, the investment decisions are dynamic and depend on the current as well as expected future states of the industry. The model allows for both product and process innovation. A higher level of knowledge may boost demand, for example by improving vehicle quality or by introducing novel product features. It may also reduce marginal cost by improving the efficiency of production. The effect of R&D investment on knowledge is the sum of a deterministic and a random component, capturing the stochastic nature of the innovation process. The technological knowledge depreciates at an exogenous rate.

There is no entry or exit in the model, as was the case in the global automobile industry over our study period (1982–2004). However, a firm may 'exit' by merging with another firm. Mergers are an important component of the model as they lead to discrete changes in market structure and force the firms to re-adjust their prices and investments to to the new industry structure. They happen for exogenous reasons, but firms take them into account when forming expectations over future valuations. In the remainder of this section we describe the demand and supply sides in some detail and then define the Markov Perfect equilibrium of the model.

3.1 The Demand Side

We motivate our demand system using a discrete choice model of individual consumer behavior, following Berry et al. [1995] and many others studying the automobile industry.[7] Vehicles differ in quality, which has two components. The first component, ω, is observable to the market participants as well as the econometrician and is a proxy for the firm's technological knowledge. The second quality component, ξ, is unobservable to the econometrician, but all firms and consumers know and use it in their pricing and purchase decisions.

There are m consumers in the market and each buys one vehicle. The utility for consumer i from buying vehicle j depends on the quality of a vehicle, its price and the consumer's idiosyncratic preferences as follows

$$u_{ij} = \theta_\omega \log(\omega_j + 1) + \theta_p \log(p_j) + \xi_j + \nu_{ij}, \quad i = 1, \ldots, m, \quad j = 1, \ldots, n. \tag{1}$$

θ_ω measures how quality conscious consumers are and θ_p determines their price elasticity. The price of vehicle j is adjusted for vehicle characteristics such as size and performance characteristics as will be described in the data section. This procedure requires the use of $\log(p_j)$ rather than the price level and results in a CES demand system.

---

[7]We did not use a random coefficient model because our main focus is on innovation and not on substitution patterns or incorporating consumer heterogeneity. A practical reason for not using the random coefficient model is to save on computation.

If we assume that the idiosyncratic utility ν_{ij} is i.i.d. extreme value distributed and there is no outside good, we obtain the following expected market share for firm j:[8]

$$\sigma_j(\omega_j, \xi_j, p_j, \mathbf{s}^{-j}, \mathbf{p}^{-j}) = \frac{\exp(\theta_\omega \log(\omega_j + 1) + \theta_p \log(p_j) + \xi_j)}{\sum\limits_{k=1}^{n} \exp(\theta_\omega \log(\omega_k + 1) + \theta_p \log(p_k) + \xi_k)}, \tag{2}$$

where $\mathbf{p}^{-j}$ is the price vector of all other firms (excluding firm j) in the industry. The expected demand is simply $m\sigma_j(\cdot)$. Each firm's demand depends on the full price vector in the industry, directly through the denominator of (2) and indirectly through its own price which, in equilibrium, is a function of its rivals' prices.

3.2 The Supply Side

We start the discussion of the supply side with the period profit function. We assume that R&D investments only generate useful knowledge with a one period lag and that prices can be adjusted flexibly period by period. After observing the industry state, firms engage in a differentiated products Bertrand-Nash game, choosing their own price to maximize profits while taking prices of rivals ($\mathbf{p}^{-j}$) and the industry state ($\mathbf{s}$) as given.

The profit maximization problem of an individual firm j is

$$\pi_j(\omega_j, \xi_j, \mathbf{s}^{-j}, \mathbf{p}^{-j}, m) = \max_{p_j}\{(p_j - \mu_j(\omega_j, \xi_j))m\sigma_j(\cdot)\}, \tag{3}$$

where μ_j is the marginal cost incurred by firm j to produce a vehicle. It is a function of the firm's knowledge, reflecting cost-reducing process innovations, and of the unobserved (to the econometrician) vehicle characteristic, reflecting cost-increasing quality improvements.

The first order condition for firm j, after some simplification, is

$$(p_j - \mu_j(\omega_j, \xi_j))[1 - \sigma_j(\cdot)]\theta_p + p_j = 0. \tag{4}$$

Since there are n firms, we have to solve n such first order conditions simultaneously to obtain the equilibrium price vector $\mathbf{p}^*$.[9] Equilibrium profits are given by

$$\pi_j(\omega_j, \xi_j, \mathbf{s}^{-j}, m) = (p_j^* - \mu_j(\omega_j, \xi_j))m\sigma_j(\omega_j, \xi_j, p_j^*, \mathbf{s}^{-j}, \mathbf{p}^{*-j}). \tag{5}$$

Fixing the functional forms and parameter values for the demand and cost functions, we can evaluate (5) for any industry state $\mathbf{s}$.

---

[8]Anderson, de Palma, and Thisse (1992, p. 136) illustrate how the CES demand system can equivalently be motivated using the representative consumer framework. Without an outside good, the discrete choice approach with unit-demand leads to the same demand equation, only the interpretation of the parameters differs.

[9]The existence and uniqueness of equilibrium in this context has been proved by Caplin and Nalebuff [1991].

The investment in R&D is a strategic and dynamic decision. Each period, firms choose their investment based on the expected value of future profit streams. The problem is recursive and can be described by the following Bellman equation

$$V_j(\omega_j, \xi_j, \mathbf{s}^{-j}, m) = \max_{x_j \in \mathbb{R}^+} \left\{ \pi_j(\omega_j, \xi_j, \mathbf{s}^{-j}, m) - cx_j + \beta E[V_j(\omega'_j, \xi'_j, \mathbf{s}'^{-j}, m')] \right\}, \qquad (6)$$

where β is the discount factor, c is the R&D cost per unit of new technological knowledge and x_j is the dynamic control, the addition to new knowledge. A prime on a variable denotes its next period value. c is the only dynamic parameter in our model that we estimate. The solution to (6) is a policy function $x_j(\omega_j, \xi_j, \mathbf{s}^{-j}, m)$.

The knowledge stock of the firm evolves as follows:

$$\omega'_j = (1 - \delta)\omega_j + x(\omega_j, \xi_j, \mathbf{s}^{-j}, m) + \epsilon^\omega_j, \qquad (7)$$

where δ is the depreciation rate of technological knowledge. ϵ^ω is a random shock that represents the uncertainty involved in doing R&D. A firm that spends cx on R&D will, on average, increase its knowledge stock by x. We assume that ϵ^ω follows a well defined and known distribution.[10]

The other two state variables are assumed to evolve exogenously and we estimate the relevant parameters over the observed period. We assume an AR(1) process with firm-specific fixed effects for the evolution of unobserved vehicle quality ξ,

$$\xi'_j - \bar{\xi}_j = \rho(\xi_j - \bar{\xi}_j) + \epsilon^\xi_j, \qquad (8)$$

where $\bar{\xi}_j$ is the long run average quality of firm j. We assume a linear trend for the size of the market m.

3.3 Mergers

Given the recent evolution of the global automotive industry, there is no need to consider entry or exit, but we do need to deal with mergers. During the 23 year period we study (1982–2004), a total of ten mergers, acquisitions and associations took place among the 23 initial firms in our sample.[11]

---

[10]This law of motion for the state variable is less general than in the main analysis in Ericson and Pakes [1995], as depreciation is deterministic, but more general than in the example they analyze in detail, as we allow x to take on a continuum of values.

[11]These were: Alfa Romeo-Fiat (1987); Ford-Jaguar (1989); GM-Saab (1990); Ford-Volvo (1998); Daimler-Benz-Chrysler (1998); Hyundai-Kia (1998); Toyota-Daihatsu (1999); Renault-Nissan (1999); Ford-Rover (2000) and GM-Daewoo (2001). In 2007 Chrysler and Daimler-Benz split; in 2008 Ford sold Jaguar and Rover brands to Tata Motors. We ignore these splits as they happened after our study period, but they could be incorporated in the same way as we treat mergers.

There is no clear pattern, making the merger activity nearly impossible to predict. Several small firms have been absorbed by larger groups, not unexpected in an industry with large scale economies. But at the same time, other small companies are thriving in spite of incessant merger predictions; Honda and BMW are consistently among the most profitable firms in the industry. Sometimes a merger is forced on a firm, as happened to Hyundai when the Korean government wanted it to bail out Kia after the Asian economic crisis in 1998. There have also been instances of two large firms joining forces in one form or another. The merger between Chrysler and Daimler-Benz is the most recent such example and the association between Nissan and Renault is another.

It is especially difficult to model merger decisions in our simple model with only two firm-level state variables.[12] Because our main interest is in the innovation decisions, we assume that mergers are random and occur with an exogenously given probability, picked to match the average observed merger rate. We do impose that the probability of merging is a declining function of n. On competition policy grounds, we further impose a lower threshold of ten on the number of active firms and we disallow mergers between firms that had a combined initial market share of over a quarter of the market.[13]

Even with this simple merger technology we need to specify the state variables for the newly created firm and the way each firm incorporates possible mergers in the evaluation of its future value. In the baseline case, we assume that the knowledge stock of the merged firm is the sum of the individual knowledge stocks of the merging firms and that the unobserved quality is the average of the vehicle qualities. We do perform a number of sensitivity checks for these crucial assumptions and estimate the model under alternative assumptions which are discussed later.

To illustrate how firms incorporate mergers in their dynamic decision making, consider a simple duopoly industry structure with firms A and B. Each period there is an exogenous probability p_m that the two firms merge. Firm A incorporates this in the calculation of its value functions as follow:

$$
\begin{aligned}
V_A(\omega_A, \xi_A, \omega_B, \xi_B) &= \max_{x \in \mathbb{R}^+} \Big\{ \pi_A(\omega_A, \xi_A, \omega_B, \xi_B) - cx(\omega_A, \xi_A, \omega_B, \xi_B) \\
&\quad + \beta \Big(p_m \zeta_A(\omega'_A, \xi'_A, \omega'_B, \xi'_B) E[V_A(\omega'_A + \omega'_B, (\xi'_A + \xi'_B)/2)] \\
&\quad + (1 - p_m) E[V_A(\omega'_A, \xi'_A, \omega'_B, \xi'_B)] \Big) \Big\},
\end{aligned}
\tag{9}
$$

where $\zeta_A(\omega'_A, \xi'_A, \omega'_B, \xi'_B)$ is the share of firm A in the total value of the merged firm. We assume

---

[12]Even small firms can translate new knowledge into higher sales instantaneously, which precludes access to the dealership system of a larger firm as a merger motive. Given that we estimate the policy function to be concave in a firm's own knowledge (see below), gaining scale economies in R&D also cannot explain mergers in our model.

[13]This is consistent with Klepper [2002], who argues that the U.S. automobile industry has settled into a stable oligopoly. We assume that the same will happen to the global automobile industry.

that it is simply the ratio of the stand-alone value of A to the sum of the values of both firms in absence of the merger:

$$\zeta_A(\omega'_A, \xi'_A, \omega'_B, \xi'_B) = \frac{E[V_A(\omega'_A, \xi'_A, \omega'_B, \xi'_B)]}{E[V_A(\omega'_A, \xi'_A, \omega'_B, \xi'_B)] + E[V_B(\omega'_B, \xi'_B, \omega'_A, \xi'_A)]}. \tag{10}$$

The approach extends to an industry with more firms, but the computations become more involved, see the Appendix for details. In particular, with more firms there are many merger possibilities and even if a particular firm is not involved in a merger itself, it has to incorporate possible mergers of its competitors in its expectations. As a result, the last two terms in (9) are replaced by a series of terms.

To summarize, in each period, the sequence of events is the following:

1. Firms observe individual and industry states.

2. They make pricing and investment decisions.

3. Profits and investment outcomes are realized and the industry state is updated.

4. A merger may (randomly) occur and the industry state is updated accordingly.

3.4 Markov Perfect Equilibrium

The Markov Perfect Equilibrium of the model consists of $V(\mathbf{s})$, $\pi(\mathbf{s})$, $x(\mathbf{s})$ and $Q(\mathbf{s}', \mathbf{s})$ such that:

1. $V(\mathbf{s})$ satisfies (17) and $x(\mathbf{s})$ is the optimal policy function;

2. $\pi(\mathbf{s})$ maximizes profits conditional on the state of the industry;

3. $Q(\mathbf{s}', \mathbf{s})$ is the transition matrix that gives the probability of state $\mathbf{s}'$ given the current state $\mathbf{s}$.

Estimating the demand and supply parameters in the period profit function is fairly standard. The one dynamic parameter in the model that we estimate, c the cost of R&D, poses a greater challenge. There are at least two ways to proceed. The first is to compute the Markov Perfect Equilibrium (MPE) from a starting value of c and use maximum likelihood estimation to fit the observed investment decisions to the predictions from the model's Euler equations. Having solved for the equilibrium, one can simulate the model and study the dynamics of interest. Benkard [2004] uses such an approach in his study of the market for wide-bodied commercial aircraft. The main disadvantage is the great computational burden, even prohibitive in our application with 23 firms, of numerically solving for the MPE over and over.

11

A number of alternative approaches have recently been developed to estimate the structural parameter(s) of a dynamic game while avoiding the calculation of the equilibrium by assuming that the data represent an MPE. As the observed investment decisions are equilibrium outcomes of the game, they can be used to characterize the policy function and the state transition matrix. This assumption is not completely innocuous. For example, firms might undergo a restructuring with major adjustments spread over several years. They might not behave as in equilibrium during such transitions. Moreover, if multiple equilibria exist, firms have to play the same equilibrium in all periods.[14]

We adopt the two-step approach proposed by Bajari et al. [2007] as it is most suitable for a continuous choice variable.[15] In the first step, state transition probabilities and the equilibrium policy functions are combined with the estimates of the period profit function to obtain numerical estimates of the value function by forward-simulation. In the second step, the minimum distance estimator chooses the dynamic parameter c that minimizes the deviations from equilibrium conditions of the model.

4   Data

The sample period 1982–2004 covers most of the consolidation that has taken place in the industry over the last few decades. The largest thirteen firms (in terms of unit sales) active in 2004 are included all the way back to 1982, as well as the ten initially independent firms they absorbed. Throughout, we do not distinguish between full and partial ownership ties, e.g. Nissan and Renault are treated as a single firm after they initiated an alliance in 1998, even though neither firm obtained majority control.

Figure 3 lists all firms in the initial and final year of the sample. In 2004, the sample accounted for more than 95% of global automobile sales. The remaining 5% of sales were made by a large number of small firms, mostly operating in only a few countries. Since patenting activity of these small firms is negligible and other information on them is spotty, we ignore them.[16]

To estimate the model we need data on the following four variables: gross additions to knowledge, knowledge stock, market share, and prices. The gross addition to a firm's knowledge is

---

[14]For a review of the problem of multiple equilibria in these models and its possible solutions see Section 6 in Doraszelski and Pakes [2007].

[15]Alternative approaches that avoid computing the MPE at each iteration include Aguirregabiria and Mira [2007], Pakes et al. [2007], and Pesendorfer and Schmidt-Dengler [2008].

[16]The only sizeable firm in 1982 not included in our sample is AvtoVAZ (Lada) in the USSR. British Leyland in the U.K. was part of the Rover Group and AMC in the U.S became 46% owned by Renault in 1979.

measured by the number of patents applied for by a firm in a calendar year,[17] taken from PAT-STAT[18]. Since firms often file for patents through various subsidiaries, we searched the database for all records containing the core of the parent firm's name and manually verified the results.

The number of applications of each firm-year observation to the U.S. and European patent offices are combined as follows: $x_{jt} = \max(x_{jt}^{US}, \lambda x_{jt}^{EU})$. λ is the relative weight given to more expensive and more demanding European patents. It is computed by taking the ratio of U.S. to European patents observed for four large firms (Daimler, Ford, Honda and Toyota) that have significant sales and production in both regions. We compute this weight to be 2.36 and conduct a sensitivity analysis with respect to this parameter.

We measure the knowledge of a firm by its patent stock. Using the number of patents applied for by each firm for the period 1982–2004, we construct the patent stock using the perpetual inventory method: $\omega_{t+1} = (1 - \delta)\omega_t + \tilde{x}_t$.[19] The series is initialized by $\omega_0 = \frac{x_0}{g+\delta}$, where x_0 is the number of new patents applied for in the first year, δ is the depreciation rate, and g is the average growth rate of new patent applications in the first five years of the sample.

Our empirical counterpart to sales is the number of vehicles sold worldwide by each firm and its affiliates. This information is obtained from Ward's Info Bank, the Ward's Automotive Yearbooks, and the online data center of Automotive News for the most recent years. The market shares are then the ratios of the number of vehicles sold by each firm to the total number of vehicles sold in the global market.

Prices for the 'composite' models are constructed by estimating a hedonic price regression for all available models in the market. The log of the list price is the dependent variable and a rich set of vehicle characteristics are the explanatory variables—see Goldberg and Verboven [2001] for an example. We include a full set of firm-year interaction dummies and the coefficients on these capture the relative price for each firm in each year (the excluded firm is GM). We estimate the hedonic regression separately for the U.S. passenger vehicle market, and jointly for five E.U. countries.[20] The weighted average of a firm's price in the two markets is the global price, normalized to zero for GM in each year.

---

[17]Patents are widely used as a measure of innovation output. In his survey on use of patents as a measure of technological progress, Griliches [1990] concludes: "In spite of all the difficulties, patent statistics remain a unique resource for the analysis of the process of technical change." [p. 1701]

[18]http://wiki.epfl.ch/patstat

[19]As mentioned before, the firm plans to obtain x_t patents, but the randomness in the innovation process yields the observed number $\tilde{x}_t = x_t + \epsilon$ of new patents, with $E(\tilde{x}_t) = x_t$.

[20]We have updated the data sets in Petrin [2002] and Goldberg and Verboven [2001] to 2004 using information from JATO Dynamics.

Figure 4 illustrates the evolution of these prices over time. After an initial increase, the price dispersion fell notably, especially from 1990 onwards. The disappearance of several independent firms with extreme prices, Jaguar and Mercedes-Benz at the high end and Kia and Skoda at the low end, is the primary explanation. Another factor was the improved reputation, and coincident price increases, for Japanese brands which caught up with and then passed GM. A second trend is the gradual decline in the average price of GM vehicles initially, which reverses at the very end of the period. The increase in the relative weight of its more upscale European brands, Opel and Saab, partly explains the higher relative GM price. Another factor is the aggressive price discounting that especially GM and other U.S.-owned brands engaged in following the 9/11 terrorist attacks. It introduced a discrepancy between the high list and lower transaction prices.

5   Estimation Methodology and Results

We now describe the estimation methodology, which closely follows Bajari et al. [2007]. In the first step, we estimate the demand and cost parameters, the policy function and transition probabilities, and use them to find numerical estimates of the value function by forward simulations. In the second step, we estimate our single structural parameter with a minimum distance estimator. We present benchmark estimates and some sensitivity checks immediately following the discussion of estimation methodology for the different elements.

5.1   Step 1

5.1.1   Demand and Cost Parameters

The demand side in our model is static and we can follow Berry [1994] to write the log of the market share of firm j relative to a base firm 0 as[21]

$$\log[\sigma_j(\cdot)/\sigma_0(\cdot)] = \theta_\omega \log[(\omega_j + 1)/(\omega_0 + 1)] + \theta_p \log[p_j/p_0] + [\xi_j - \xi_0]. \tag{11}$$

Using the observed market shares and data on the knowledge stock and prices, we can estimate the above equation pooling all our data. To obtain consistent estimates, we need to account for the endogeneity of prices. As firms use information about the unobserved vehicle quality in their pricing decisions, prices will be correlated with the error term and OLS estimates will be inconsistent. In particular, we expect the price coefficient to be biased upward. We follow the identification strategy

---

[21]We estimate demand, cost, and policy functions pooling data across all years; time subscripts are omitted. Throughout, we use GM as the base firm.

of Berry et al. [1995] and use the sum of observable characteristics of rival products as instrument. In our case, this boils down to the sum of rivals' knowledge.

A second complication is the assumed AR(1) process with fixed effects for the error term ξ_j, as introduced in equation (8). After de-meaning all variables to take care of the fixed effects, we follow a two-step procedure. First, the persistence of the shocks to ξ is estimated as the coefficient on lagged market shares in the demand equation. Second, we quasi-difference the dependent and explanatory variables in the demand equation and estimate using instrumental variables.[22]

To characterize the dependence of marginal costs on the state variables, the two quality components, we estimate the first order condition (4) jointly with the demand equation. In the benchmark case, we use a quadratic specification for the relative marginal cost function:

$$\mu_j/\mu_0 = \gamma_0 + \gamma_1(\omega_j/\omega_0) + \gamma_2(\omega_j/\omega_0)^2 + \gamma_3\tilde{\xi}_j + \gamma_4\tilde{\xi}_j^2 + \epsilon_j^\mu \tag{12}$$

with $\tilde{\xi}_j = \exp(\xi_j - \xi_0)$.

The estimates for the demand and cost parameters in the benchmark case are reported in Table 2. Price enters negatively and with our CES demand specification the coefficient has the interpretation of the price elasticity of demand. At –2.35, demand is estimated to be less elastic than in studies using micro-level data (see, for example, Berry et al. [1995]). This is consistent with the more aggregated definition of a product. The residual demand for a firm will be less elastic than for an individual model. The implied price-marginal cost markups are correspondingly lower. Again, this is reasonable because at the firm level, as opposed to the model level, a larger fraction of costs will be variable.

Estimation demand jointly with the first order conditions for optimal price setting guarantees an elasticity above unity, but our estimate is significantly different from this bound.[23] Weak instruments or controlling imperfectly for the product characteristics in the hedonic price regression would lead to an upwardly biased price coefficient. Both issues do not appear to be problematic.

The impact of knowledge, proxied by the firms' patent stock, on sales is estimated positively and precisely. This variable helps to control for some of the unobserved quality variation that can lead to an upward bias and even a positive estimate for the price coefficient in micro-level studies. A one percent increase in relative knowledge stock is estimated to boost relative market share by 0.17 percent. A firm with a patent stock half of GM's, will on the average have a market share that is 12% ($= 0.17 \times \ln(0.5)$) lower than GM's, holding prices constant.

---

[22]A robustness check using the GMM-SYS estimator yielded similar results, but the system estimator is difficult to combine with the joint estimation of the first order condition.

[23]Estimation of the demand system by itself, without enforcing the first order conditions for optimal price setting, always produced estimates for demand elasticity higher than unity.

Greater knowledge is estimated to have a concave effect on marginal costs. The linear effect is positive. To incorporate the benefits of increased knowledge, it is often necessary to introduce new features into vehicles, which is costly. The negative coefficient on the squared knowledge term indicates that the marginal effect on costs decreases as knowledge grows larger. Patents for process innovations are one possible explanation for the negative coefficient, if they are more prevalent for highly innovative firms.

The negative relationship between marginal costs and the unobserved quality term (ξ) is less intuitive. It suggests that firms are able to produce higher quality vehicles at lower costs. The point estimate on the squared term comes in positive, but it is not significantly different from zero and of very small magnitude. This quality dimension could capture an attractive design or another feature that might require high fixed costs, but does not increase the marginal production costs.

The estimates in Table 2 will be used in the benchmark estimation of the dynamic parameter c. A number of sensitivity checks will be performed using alternative estimators or specifications for demand and costs. These results are summarized in Table 3.

The first column reports estimates where no instruments are used for price in the demand equation. It leads to a less elastic demand and a higher coefficient for knowledge. Both effects are as expected because firms will price quality vehicles higher. Note that this change also has effects on the cost parameters. All coefficients retain their signs, but the absolute magnitudes of the parameters on the knowledge variables are lower and those on the ξ variables are higher.

In column (2), the unobserved vehicle quality is modeled as constant over time. It eliminates the AR(1) process for the error term in the demand equation and allows for a much simpler estimation procedure. The impact of knowledge on market share is estimated to be a lot stronger, with little change in the demand elasticity. In the marginal cost equation, the impact of quality is now solely identified from variation across firms, eliminating variation over time. It enhances the estimated negative effect of ξ on marginal costs, but also leads to a negative effect of knowledge on costs. The high observed correlation between ξ and ω, 0.71 in the benchmark case, explains the spillover effects on the parameters on the knowledge variables, as these now pick up all variation over time.

In columns (3) and (4) the functional form of the marginal cost equation is varied. A linear specification highlights the positive effect of ω and the negative effect of ξ. A log-linear specification does the same, only the effects will be reduced at high levels of knowledge or quality. In both cases, the demand equation is estimated to be slightly more elastic, but the changes are minor.

5.1.2 Policy Function and State Transitions

To calculate future profits, we need to calculate future values for all state variables. As it is computationally infeasible to estimate the optimal policy function for investments, we follow Bajari et al. [2007] and characterize an empirical policy function using observed investment decisions. Ideally, the control variable x should be modeled as a completely flexible function of a firm's own state and the full vector of its rivals' states that are contained in the industry state. Bajari et al. [2009] discuss the resulting distribution of the estimator if this step is carried out nonparametrically.

Given the data limitations, we use the following specification:

$$x_j = \sum_{k=0}^{3} \sum_{l=0}^{3-k} \sum_{m=0}^{3-k-l} \alpha_{klm}(\omega_j)^k (\omega_{-j})^l (\xi_j)^m + e_j, \tag{13}$$

where $\omega_{-j} = \sum_{i \neq j} \omega_i$, the total patent stock of firm j's competitors. e_j is an approximation error between the innovation that the true policy function would generate and our prediction.

To avoid extreme values for the number of new patents when we use equation (13) to make far out of sample predictions, the parameters on terms with third powers are set to zero. While this is a restrictive specification compared to the optimal investment policy, which is an unknown function of 46 continuously valued variables (two state variables per firm), the fit of the regression is very good with an R-squared of 0.938. When the equation is used to predict future innovations, the values are restricted to be weakly positive.

The full set of coefficient estimates are reported in column (1) of Table 5 in the Appendix. We also report results for a restricted policy function with fewer terms, in column (2), and for a policy function where the logarithm of the state variables instead of the levels enter the polynomials in equation (13). These alternative specifications will be used to verify the sensitivity of our results. Note that a different policy function is also estimated (but not reported) when the alternative estimates for the demand and cost parameters from Table 3 are used.

The state transition function for the knowledge stock ω is then given by equation (7). The accumulated stock of past patents decreases exogenously in value with depreciation and increases with new innovations. We assume a depreciation rate δ of 15%, the same value used in the construction of the patent stock series, taken from Cockburn and Griliches [1988]. It captures both patent expirations and the economic obsolescence of older knowledge. In the robustness checks we also estimate the dynamic parameter assuming alternative depreciation rates of 10% and 20%.

The stock of knowledge increases with the number of newly granted patents, which differs from the desired number of patent applications in equation (13) due to the randomness in the discovery and patent granting process. The error ϵ^ω is assumed to follow a normal distribution with mean

zero and standard deviation that is 10% of x. The addition to the knowledge stock will thus be $x_j(1+\epsilon^\omega)$. A firm that aims to apply for 100 new patents in a particular year has a 67% probability of ending up with 90 to 110 new patents.[24]

The unobserved quality terms ξ are recovered as residuals from the demand equation. From the same estimation we also recover the ρ parameter in the AR(1) process for ξ. Equation (8) provides an estimate of the standard deviation for the shocks to ξ, which is estimated relatively small (standard deviation of 0.21) compared to the variation in ξ across firms (standard deviation of 0.92). Shocks to ξ are however very persistent, with a decay parameter of 0.81. In the robustness checks, we verify the impact of eliminating the time-variation in ξ.

An important feature of the estimated policy and transition functions is that when the industry evolution is simulated forward, the estimated number of new patent applications does not increase without bound. After some transition period, the industry settles in a (stochastic) steady state, where firms' stock of patents hoover around a certain value, which will differ across firms. The actual number of new patent applications in each year will vary with the random shocks to the innovation process and to the transition process for ξ.

The final ingredient in the computation of future profits is a prediction of the future size of the industry. We estimate a linear trend for (global) industry sales over our sample period and find an annual addition of 781,000 vehicles per year.

5.1.3 Computation of the Value Function

Combining the previous building blocks with a vector of dynamic parameters θ, we can calculate an estimate for the value function from any state $(\omega_j, \xi_j, \mathbf{s}^{-j}, m)$. In our model, θ consists of a single parameter: c, the R&D cost required to obtain one new patent in expectation. For now, the evaluation of the value function will be conditional on a starting value c_0. In step 2 below, we use the equilibrium conditions of the model in a minimum distance estimator to estimate this parameter.

To evaluate the value function we use forward simulations. We first explain this procedure in the case without mergers. For an initial industry state $\mathbf{s}_0$, the estimated demand and cost equations allow the calculation of the equilibrium price vector from the system of first-order conditions (4). The estimated policy function directly determines optimal innovation. Substituting these in the

---

[24]An alternative approach would be to define the control x as the value of R&D expenditures. We could then estimate a patent production function $\Delta\omega_t = f(\omega_{t-1}, x, \epsilon)$ from the data, and let knowledge evolve according to $\omega_t = (1-\delta)\omega_{t-1} + \Delta\omega_t$. In our approach, the cost of innovation is estimated at cx and we will compare this with observed R&D expenditures as a reality check.

equilibrium period profit (5) for all active firms gives the net profit vector: $\pi_0(\mathbf{s}_0) - c_0 x_0(\mathbf{s}_0)$.

Next, we use the state transition functions to find the industry state for the next period $\mathbf{s}_1$. We draw random errors, ϵ^ω and ϵ^ξ, for each active firm, which determine the future states according to equations (7) and (8). Following the same steps as above we can compute the expected net profit in period 1 as $\pi_1(\mathbf{s}_1) - c_0 x_1(\mathbf{s}_1)$, which is discounted back to period 0 using a discount factor β. The benchmark value for β is 0.92, but we later report sensitivity analysis with respect to this parameter.

This process is repeated T periods, until β^T gets arbitrarily close to zero. In the benchmark estimation we set T to 150 periods and construct the value functions as the present discounted value of the profit streams.[25]

In sum, we evaluate the following equation:

$$V(\mathbf{s}_0|c_0) = \mathbb{E}\left[\sum_{t=0}^{T} \beta^t [(\pi_t(\mathbf{s}_t) - c_0 x_t(\mathbf{s}_t)]\right], \tag{14}$$

where the expectation is over future states. We perform these forward simulations many times, using different draws on the error terms (ϵ^ω and ϵ^ξ), the only sources of uncertainty in the model. The average is the numerical estimate of $V(\mathbf{s}_0|\theta)$.

This algorithm becomes more complicated in the presence of mergers. At the end of each period, all firms receive a draw from the uniform distribution over the unit interval. If in some period two firms receive a draw below p, they merge. The value for p is calculated in the Appendix to match the observed merger frequency p_m, which was 10/23 over the sample period, and to generate a declining probability of merging if the number of active firms shrinks.

After a merger, the ω and ξ of the merged firm are, respectively, the sum and the average of the values for the individual firms. This changes the state of the industry in all calculations from that point onwards. Below, we report robustness checks using alternative assumptions to construct the state values for the newly merged firm.

To divide the future profits of the merged firm to the value functions of each of the merging firms, we use the share of each firm in the combined value should they have remained independent—according to equation (10). This still requires the calculation of future patent stocks for all firms and the profits for the two merging firms as if the merger had not taken place in all future periods. It proved computationally infeasible to perform these calculations allowing for further mergers in the future. Instead, the ζ terms are calculated using value function realizations in the absence of any mergers.

---

[25]The estimate for c is virtually unchanged if T is increased beyond 150, but the computation burden rises proportionally.

The randomly occurring mergers add an important new source of randomness to the forward simulation of values. The estimate for c only became invariant to starting values under two conditions. We increased the number of error sequences (including mergers) that we average over in the calculation of the expected value function in equation (14) to 300. To avoid excessively concentrated market structures, mergers are only allowed if the sum of the initial market shares of the merging firms did not exceed one quarter and if at least ten independent firms remain.

5.2 Step 2

In step 2 we use the results from the first stage together with the equilibrium conditions on the MPE to recover the dynamic parameter of the model. It assumes that the model is identified and there is a unique true parameter vector θ^*. Bajari et al. [2007] propose a minimum distance estimator for this true parameter vector. Let $\mathbf{x}(\mathbf{s})$ be the equilibrium policy profile. For this to be a MPE policy profile, it must be true that for all firms, all states, and all alternative policy profiles $\mathbf{x}'(\mathbf{s})$

$$V_j(\mathbf{s}, \mathbf{x}(\mathbf{s}), \theta) \geq V_j(\mathbf{s}, \mathbf{x}'(\mathbf{s}), \theta), \tag{15}$$

where $\mathbf{x}' \neq \mathbf{x}$ only at the jth element, while all other firms play their Nash strategy. Equation (15) will hold at the true value of the parameter vector θ^*.

The minimum distance estimator for θ^* is constructed as follows. For each firm j and each state $\mathbf{s}$ observed in the sample, we use the forward simulation algorithm to calculate $V_j(\mathbf{s}, \mathbf{x}(\mathbf{s}), \theta)$. We do the same calculations for a number of alternative policy profiles $\mathbf{x}'(\mathbf{s})$ and compute the difference $d(j, \mathbf{s}, \mathbf{x}'|\theta) = V_j(\mathbf{s}, \mathbf{x}(\mathbf{s}), \theta) - V_j(\mathbf{s}, \mathbf{x}'(\mathbf{s}), \theta)$. For a given value of θ, we calculate these differences for all j, $\mathbf{s}$ and $\mathbf{x}'(\mathbf{s})$ and compute the sum of the squared $\min\{d(j, \mathbf{s}, \mathbf{x}'|\theta), 0\}$ terms. Note that we only consider the negative differences because they imply a violation of equilibrium condition (15). θ is chosen to minimize these violations, i.e.

$$\hat{\theta}^* = \arg\min \sum_{j, \mathbf{s}, \mathbf{x}'} \left[\min\{d(j, \mathbf{s}, \mathbf{x}'|\theta), 0\} \right]^2. \tag{16}$$

To evaluate the objective function in the minimum distance estimator, we have to specify alternative policies to the equilibrium Nash policy. For each firm j and each industry state $\mathbf{s}$ we specify two alternative policies as: $\mathbf{x}'(\mathbf{s}) = (\iota \pm e_j)'\mathbf{x}(\mathbf{s})$, where ι is a vector of ones and e_j is a vector of zeroes with 0.01 at position j (both vectors are of length n). In words, the alternative policies are to choose x which is one percent less or one percent more than the equilibrium policy.

The only dynamic parameter in our model is the (average) R&D cost of obtaining one new patent (c). The benchmark estimate, using preferred values for all the static parameters and preferred

specifications for all primitives as discussed above, is an estimate for $\hat{c}$ of 21.8 million dollars per targeted new patent application.[26] Using subsampling with 12 years of data and 20 iterations, we obtain a standard error of 10.55 million dollars.[27]

5.3 Sensitivity analysis

A first reality check for the \$21.8 million benchmark estimate for the cost of innovation parameter c is to compare it with observed R&D expenditures per patent application. For a subset of firms in the later years of our sample we can calculate this ratio as it has become compulsory for U.S. firms to report R&D expenditures on their annual 10-K form. With an observed average of \$15.6 million and a median of \$14.9 million, the R&D data seems broadly consistent with our estimate.

We also verified how robust the estimate is to different assumptions on functional forms for various primitives, estimation procedures for static parameters, and other parameter choices. Results on all sensitivity checks are collected in Table 4. In each panel (a) to (g), one ingredient is changed at a time. The assumption under the benchmark estimation is always listed on the first line, followed by the benchmark estimate. In subsequent rows, the assumption is changed and the new c estimate is reported.

Results in panel (a) of Table 4 illustrate that estimating the demand system with OLS or fixing the unobserved product quality to a constant would raise the estimate. In the first case, the driving force is the much lower demand elasticity estimate, see column (1) of Table 3. As a result, firms can charge higher markups and a knowledge stock advantage becomes a lot more valuable. In order to fit the observed rate of patenting, the estimator assigns a higher R&D cost to each patent. In the second case, fixing ξ to a firm-specific constant more than doubles the coefficient on knowledge in the demand equation. Again, to fit the observed patenting rate in an environment where patents are more valuable the estimated cost of patenting rises.[28]

For the results in panel (b), we revert to the benchmark demand model. Here, the marginal cost equation is changed instead. We report results for two linear specifications using the logarithms or the levels of the state variables as arguments. Using logarithms, see column (3) in Table 3, the basic shape of the quadratic specification can be maintained and costs are estimated to increase concavely in ω and decrease convexly in ξ. It is no surprise then that the c estimate changes only by 21%—about two-fifth of a standard deviation. With only the linear effects of the levels included

---

[26]This number is based on a normalized (real) price of \$10,000 for a GM vehicle.

[27]The high standard error is largely due to the re-estimation of demand, cost and policy functions for each sub-sample.

[28]Other (unreported) robustness checks confirm this mechanism. For example, including time dummies in the demand equation raised the price elasticity estimate and the lower optimal mark-ups translates into a lower c estimate.

in marginal costs, the difference is even smaller, with a c estimate of \$22.7 million.

In the Bajari et al. [2007] estimator for dynamic games, the policy function that characterizes the observed decisions plays a key role. Our preferred specification, in equation (13), is non-parametric in levels. We also tried a more parsimonious policy function, omitting all third order and some interaction terms, and a third fully flexible function with all explanatory variables expressed in logarithms. Coefficient estimates for all functions are reported in Table 5 in the Appendix.

The restricted specification produced a policy function more strongly concave both in the stock of own patents and in the stock of competitors' patents. This had two effects. First, patenting rates were predicted to be somewhat higher initially, leading to more rapid convergence of the industry to the steady state. Second, in the steady state which is characterized by high average knowledge stocks, stronger concavity leads to lower average patenting rates. Both changes contributed to the significantly lower estimate for the R&D cost of \$14.3 million.

It is difficult to compare the policy function with logarithms directly with the benchmark specification. The signs on many polynomial terms switched and the magnitudes of the coefficients are entirely incomparable. Expressing all knowledge stocks in logarithms before entering them in the polynomials did reduce the estimated concavity of the policy equation, with opposite effects from the changes with the restricted function. The estimate for c barely changed (dropped to \$21.0m).

In the benchmark model we assumed that when two firms merge, their combined knowledge stock is the sum of their individual pre-merger knowledge stocks and the new unobserved product quality is the average of the individual firm qualities. The results in panel (d) retain the first assumption, but change the second. Under both alternative modeling assumptions—using the highest ξ for the merged firm or taking a knowledge-weighted average—the value of a merger increases. This provides a first boost to innovation incentives for firms as long as any future merger is possible. Given that both ω and ξ enter the random utility in logs, the return to knowledge is increasing in ξ. The higher ξ's thus provide a second boost to innovation incentives even when no further mergers are possible.

The results bear out these effects as both changes lead to higher estimates for the c parameter. Picking the maximum ξ increases the estimate for c by 32%, whereas merely weighting the ξ's raises the estimate by 24%. To some extent, these are overestimates due to the constancy of the policy equation. Innovation is estimated to be increasing in own-ξ, which is intuitive. However, this effect is assumed to remain equally strong even as we calculate ξ's to be higher after mergers in the alternative estimations.[29]

---

[29]When two firms merge during the sample period, we can estimate the new ξ of the merged firm and we have

Next, we investigate the sensitivity of the c estimate to three parameters that were fixed exogenously. Raising the discount factor β from 0.92 to 0.94 or lowering it to 0.90 has barely any effect. We find the estimates in panel (e) to be slightly increasing in β. A higher value for the discount factor places more weight on the future, which is when the value of patents will be realized. It is not surprising that it raises the estimate of c, albeit only slightly from \$21.8 to \$23.8 million.

In the calculation of the patent stock data and in the forward simulations, we used a depreciation rate of 15%. A lower depreciation rate corresponds to a longer useful life span for inventions, enhancing innovation incentives, and *ceteris paribus* this should lead to a higher c estimate to match the observed patenting rates. However, the estimates in panel (f) of Table 4, suggests an opposite effect. The estimate for c, the cost of R&D, is increasing in the depreciation rate.[30]

The reason is that strategic interaction in the innovation decisions cannot be ignored. The direct effect of lower depreciation is to boost innovation, but this affects all firms similarly. As a result, firms will more quickly reach an area of the policy function where the large size of their competitors' patent stocks diminish own patent applications. Note that six of the nine terms in the policy polynomial that include competitors' patent stock are estimated to have a negative coefficient in column (1) of Table 5.[31] As a result, the net effect of slower depreciation is to reduce steady state innovation and to lower the c estimate. A reduction in the depreciation rate by a third, from 0.15 to 0.10, lowers the estimate from \$21.8 to \$18.1 million.

When calculating knowledge stocks from observed patent applications, we treat European and American patents asymmetrically. The application cost is higher and the success rate lower in the E.U., leading to fewer but on average more valuable patents. Based on observed patenting rates of firms with similar activities on both continents, we imputed a relative value of 2.36 for a European versus an American patent application. If we use a ratio that is only half as large, the estimate for c increases to \$26.8m. On the other hand, if we boost the relative value of a European patent by another 50% to 3.54, we find a much lower value for c of only \$13.0 million. The mechanism is similar to that for a lower depreciation rate: it enhances the negative effect of competitors' knowledge stock in the policy function.

---

used those in the policy function. An unweighted or sales-weighted average of the pre-merger ξ's provided the best prediction, but with only 10 mergers in the sample period, this is necessarily an imprecise assessment.

[30]Note that we also recalculate the patent stocks in the data for the alternative estimations.

[31]In particular, the term on $\omega_j(\omega_{-j})^2$ is one of the most precisely estimated and it is negative, just as the uninteracted linear and quadratic terms of $(\omega_{-j})^2$.

6 Effect of market structure on innovation

To study the effect of market structure o innovation using the estimated model, we need to compute or approximate the Markov perfect equilibrium strategies of the firms for all possible configurations in the state space. We can then use these strategies to forward simulate the industry from any initial state to study the equilibrium dynamics.

6.1 Naive Strategies

Computing the Markov Perfect Equilibrium (MPE) of a dynamic game with more than twenty players, two state variables and the possibility of mergers is not feasible.[32] An alternative would be to compute the Oblivious Equilibrium (OE) of the game (Weintraub et al. [2008]). Here, each firm assumes the industry state is fixed at some long run steady state and chooses its strategy accordingly. In a rational expectation equilibrium, the strategies chosen by the firms will lead to the same long run state of the industry that each firm took as given. The computation of the OE involves finding the terminal industry state and then solving a single-agent dynamic programming problem.

There are two important reasons we do not follow this approach. First, OE is a good approximation to MPE only when the number of firms in the industry is large and the long run distribution of the industry satisfies a thin-tail condition. In our case, the number of active firms ranges from only 10 to 23 per year and the four largest firms control almost half of the market. Second, our model features exogenous mergers until a lower bound of ten firms is reached. In this process, the unobserved quality differences are updated by averaging across merging firms. By construction, the eventual steady state will be a concentrated industry with ten firms of comparable size and knowledge stocks. The optimal innovation strategy for firms in initial situations with many more very heterogeneous firms is likely to differ a lot from the steady state innovation decision of a single agent optimizing against a constant industry state. In addition, because of the need to average over different random merger sequences, the computation of the industry steady state would also be a lot more computationally demanding.

In the absence of a feasible or reasonable alternative, we introduce what we call *naive strategies* as approximations to MPE strategies. We simply impose that each firm chooses its level of innovation by assuming that the states of its rivals will remain the same. Clearly, if the industry is not in

---

[32]The computation suffers from the so called 'curse of dimensionality' and the size of the state space as well as the computation time needed for one iteration in computation of equilibrium increases exponentially with the number of players in the game. This is a well known problem in this literature. For an detailed discussion of the computational issues in these models see Doraszelski and Pakes [2007].

a steady state, these naive strategies are not equilibrium strategies.[33] However, in our application they are likely to provide better approximations to MPE strategies than the strategies based on OE. Tomorrow's expected state (which we cannot compute) is more likely to be closer to today's actual state than to the long run industry state. Once the industry is in a steady state, both naive and OE strategies will converge to MPE strategies.

On the one hand, naive strategies are easier to compute than OE strategies because it is not necessary to find the terminal state of the industry. On the other hand, for each new state of the industry that materializes, one needs to calculate the new naive strategies appropriate for that that state, which adds to the computational burden. The forward simulations reported below took one week of CPU time to complete.

Before simulating the full-fledged model using naive strategies, we compare in a highly simplified version of the game the simulated industry dynamics for the naive and the optimal MPE strategies. In particular, we limit the industry to five firms and set the values of ξ equal to zero for all. We discretize the state space for the knowledge stock to just twenty possible values, ranging from 300 to 2,200 in increments of one hundred. We compute the two different strategies of this game, and simulate the industry forward from the following initial state $\mathbf{s}_0 = [2200\ 1800\ 1400\ 1000\ 600]$.

The evolution of knowledge for each of the five firms under MPE strategies is shown in the top panel of Figure 5. The comparable evolution for the naive strategies is calculated iteratively, updating the industry state after each period, and depicted in the bottom panel.[34] In both cases the industry converges to a steady state in about 12 periods. Without intrinsic firm heterogeneity—firms share the same ξ—all firms converge to the same average knowledge stock. The coarse state space required to compute MPE strategies leads to an identical stock of 600 patents for all firms in that case. Under the naive strategies, the randomness in the innovation process leads to some residual variation around the 600 mark. The evolution and relative knowledge positions of the different firms are very similar under the two scenarios. For our simple model of innovation, naive strategies generate industry dynamics that are similar at least in some respects to those generated by MPE strategies.

---

[33]A simple equilibrium strategy in this case would be for each firm to do just enough innovation to replace its depreciated knowledge stock.

[34]In the latter case, we can allow a much finer state space and the knowledge stocks can take any integer value from 1 to 2200.

6.2 Market structure and innovation: Forward simulation

We now return to the basic question motivating this study: what is the equilibrium relationship between market structure and innovation? To this end, we simulate the estimated model forward starting from the actual state of the industry in 1982. This will provide information how innovation is predicted to evolve as the industry converges to a steady state.

We initiate the model with the actual ω's and the estimated ξ's for the 23 active firms. We solve for the equilibrium price vector in the static pricing game using the estimated demand and cost parameters. We assume that firms follow the naive strategies to determine innovation in new knowledge. Randomly drawn shocks to innovation and unobserved quality are added when we update all state variables. If a merger takes place, which is also drawn randomly, the affected firms' states are updated further. This completes the simulation for one period. We repeat the above sequence for fifty periods and record a number of statistics after each period. Mergers are allowed until the number of active firms is down to ten.

In Figure 6 we show the evolution of all knowledge stocks. Mergers appear as an up-tick in the stock for one firm and the disappearance of the stock for another firm. For the particular merger sequence reported, the industry reaches its minimum of ten firms after the 28th period. Apart from mergers, equilibrium dynamics suggest a very rapid convergence to a stochastic steady state, which is completed about ten years after the last merger. From that point onwards, the 10 remaining firms still differ in long-term ξ's, leading to small average knowledge stock differences. However, the persistent differences are swamped by the effects of the random shocks to ω and ξ. The ω's keep increasing over time because of the growing market size.

We now illustrate how this evolution is partially shaped by and feeds back into firms' optimal investment policies. The relationship between competition and innovation intensity at the firm level is depicted in Figure 7. Competition is measured by one minus the Lerner's index, increasing to the right on the horizontal axis; innovation is measured by the ratio of R&D expenditure to total revenue. While there is a lot of variation in innovation intensity for any given level of competition, especially at low levels of competition, there is a discernable negative relationship. On average, the equilibrium innovation intensity falls with a decline in price-cost margins. While this is consistent with the classic Schumpeterian argument that market power is a pre-requisite for innovative activity, there are other drivers in the model.

High levels of innovative activity tend to corresponds to periods when firms are trying to catch up with a rival. This happens, for example, after a merger which gives one firm a large knowledge advantage. In the absence of any adjustment costs, rivals will innovate a lot to catch up with the

leader right away. Equilibrium dynamics also suggest that there is no real incentive for the leader to innovate to escape competition. Because of our symmetric demand model with concave effects of knowledge on consumer utility, the marginal benefit of innovation will fall below the (constant) marginal cost first for the knowledge leader. Persistent ξ differences modify this tendency somewhat, but in the merger process remaining firms become gradually more alike.

Another feature in the figure are the very low levels of innovation at the lowest observed levels of competition. This corresponds to the periods when the current knowledge stocks of many firms are above the steady-state knowledge. These firms basically stop innovating temporarily and let their knowledge depreciate to the steady-state value. That also happens to most newly formed firms for a couple of periods after a merger.

In Figure 8 we show the relationship between the same two variables, but averaged over all firms active in a given year. The numbers next to each point indicate the year into the simulation, starting with 0 in 1982.

As the industry was converging to the stochastic steady state—in the first 28 points but ignoring the years with intense merger activity (20, 23, 25, 26)—we observe a similar negative relationship as in Figure 7. Average markups decline as the industry consolidates. This might appear counter-intuitive, but it reflects the greater symmetry of surviving firms as the industry consolidates to the steady state. Outliers with high ξ's are absorbed in larger groups and firms with excessive ω's let their knowledge depreciate. Although the number of competitors becomes smaller, they become more evenly matched in both ξ and ω and the intensity of competition increases. This is one example where a fully specified dynamic model can generate unusual, but reasonable predictions.

Once the industry has consolidated and reached the steady state, there is no further change in the average level of competition, although there is still some heterogeneity at the firm level. The average level of innovation serves to replace depreciated knowledge and to gradually boost knowledge stocks as the market keeps expanding. The relationship between competition and innovation at the industry level disappears once the steady state is achieved.

It is instructive to look a the relationship between innovation and market structure using an alternative definition of competition. In Figure 9 we use one minus the Herfindahl index on the horizontal axis. This measure can even be calculated on the raw data, without a need for estimated price-cost margins. The degree of competition is now estimated to decrease continuously as the industry consolidates. The accompanying increase in concentration leads to a continuous leftward shift on the graph. Again ignoring the merger episodes (20, 23, 25 and 26), the relationship between competition and innovation is now mostly *positive* over the consolidation episode. The leftward shift on the graph is accompanied by declining rates of innovation. Once the consolidation is completed,

from period 29 onwards, there is again no clear relationship between competition and innovation anymore.

7   Concluding Remarks

We construct a dynamic game-theoretic model of the global automobile industry that incorporates forward-looking innovation and random mergers. We estimate the structural parameter of the model using a new method proposed by Bajari et al. [2007]. Our application illustrates the usefulness of their method to obtain parameter estimates where the computational burden makes calculation of the equilibrium impossible and we highlight some of the problems we faced when implementing their methodology.

After estimating all structural parameters, we use forward simulations to study the relationship between market structure and innovation. Our findings can be summarized as follows:

1. At the firm level, there is a weak association between lower markups and lower innovation intensity. This implies a negative relationship between competition and innovation.

2. At the industry level, our model converges to a constant market structure with only replacement innovation. During the transition period, industry consolidation is accompanied with declining innovation intensity while markups decrease. Hence in our simulations the two commonly used measures of competition, price-cost markups and market share concentration, are negatively correlated. The relationship between competition and innovation depends on the preferred definition.

3. The convergence to the steady state makes a firm's relative position in terms of technological knowledge a key determinant of its innovation intensity. In equilibrium, firms with higher knowledge relative to their rivals tend to innovate less.

The trade-off between *replacement* and *efficiency* effects that has featured prominently in the literature does not appear here. In our model, new innovations add to the existing knowledge without replacing it, eliminating the replacement effect. Existing knowledge depreciates at a fixed rate, which is independent of a firm's innovation decision. The efficiency effects are present— lower innovation rates coincide with lower markups—but they can easily be dominated by strategic considerations. For example, after transformative mergers in periods 20, 23 25 and 26, the remaining firms are faced with a much larger rival and they immediately boost innovation to catch up with the leader (even if markups are low).

A number of important features of industry dynamics did not enter our model, simply because they are less important in the automotive industry. The market size grows at a constant rate over time, typical of very mature industries. No new firms enter the market, which was the case over the 1982-2004 sample period. Some emerging Chinese and Indian producers should be included past our sample period. We incorporated mergers are exogenous events, due to their low frequency and unpredictability in the automotive industry. Our modeling environment has the potential to endogenize mergers, and we are exploring this in an application to another industry.

The absence of adjustment costs to the state variable, technological knowledge, generates unrealistic dynamics. All firms can now adjust their knowledge stocks to a new optimal equilibrium value in just one period. The only constraint is that the gross change in knowledge cannot be negative, which limits the downward adjustments to the rate of depreciation. Introducing adjustment costs is conceptually straightforward, but it would increase the computational burden tremendously as a second dynamic parameter would appear. We consider this to be an interesting and challenging topic for future research.

The benchmark model has another unrealistic implication: mergers are nearly always value destroying. The reason is twofold: the decreasing marginal returns to knowledge and the loss of an independent firm. Remedying the first effect fully would require a drastically different utility specification. One alternative treatment of mergers that we used in the sensitivity checks, assigning the maximum ξ to the merged entity, does allow for value-creating mergers by leveraging the higher knowledge stock with a higher ξ. The second effect could be remedied by leaving the number of products in the market unchanged after a merger. The new firm would simply determine prices and innovation jointly for both products, while sharing the full knowledge stock across all products.

The above limitations notwithstanding, we believe that this study is a step forward in the literature on the interaction between market structure and innovation. In particular, we have endogenized the market structure by allowing two-way interactions between innovation and firms' competitive position. The random mergers introduce further exogenous and discrete changes in the market structure, to which all firms subsequently adjust.

In line with earlier surveys, we believe that the relationship is industry-specific. While generalizations across sectors are likely to be of limited policy relevance, the framework employed in this paper provides a flexible approach to study oligopolistic-dynamic interactions in other industries.

References

P. Aghion, N. Bloom, R. Blundell, R. Griffith, and P. Howitt. Competition and Innovation: An Inverted-U Relationship. *Quarterly Journal of Economics*, 120(2):701–728, 2005.

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

S. P. Anderson, A. de Palma, and J.-F. Thisse. *Discrete Choice Theory of Product Differentiation*. The MIT Press, Cambridge, MA, 1992.

P. Bajari, C. L. Benkard, and J. Levin. Estimating Dynamic Models of Imperfect Competition. *Econometrica*, 75(5):1331–1370, September 2007.

P. Bajari, V. Chernozhukov, H. Hong, and D. Nekipelov. Nonparametric and Semiparametric Analysis of a Dynamic Discrete Game. Working Paper, April, 2009.

C. L. Benkard. A Dynamic Analysis of the Market for Wide-Bodied Commercial Aircraft. *Review of Economic Studies*, 71(3):581–611, 2004.

S. Berry. Estimating Discrete-Choice Models of Product Differentiation. *Rand Journal of Economics*, 25(2):242–262, 1994.

S. Berry, J. Levinsohn, and A. Pakes. Automobile Prices in Market Equilibrium. *Econometrica*, 63 (4):841–890, 1995.

R. Blundell, R. Griffith, and J. V. Reenen. Market Share, Market Value and Innovation in a Panel of British Manufacturing Firms. *Review of Economic Studies*, 66(3):529–554, 1999.

J. Boone. Competitive Pressure: The Effects on Investments in Product and Process Innovation. *RAND Journal of Economics*, 31(3):549–569, 2000a.

J. Boone. Intensity of Competition and the Incentive to Innovate. *International Journal of Industrial Organization*, 19:705–726, 2000b.

A. Caplin and B. Nalebuff. Aggregation and Imperfect Competition: On the Existence of Equilibrium. *Econometrica*, 59(1):26–61, 1991.

W. Carlin, M. Schaffer, and P. Seabright. A Minimum of Rivalry: Evidence from Transition Economies on the Importance of Competition for Innovation and Growth. *Contributions to Economic Analysis and Policy*, 3(1):1–43, 2004.

A. T. Ching. Consumer Learning and Heterogeneity: Dynamics of Demand for Prescription Drugs After Patent Expiration. Working paper, University of Toronto, 2008.

I. Cockburn and Z. Griliches. Industry Effects and Approbriability Measures in the Stock Market's Valuation of R&D and Patents. *American Economic Review*, 78(2):419–423, 1988.

W. M. Cohen and R. C. Levin. Empirical Studies of Innovation and Market Structure. In R. Schmalensee and R. D. Willig, editors, *Handbook of Industrial Organization*, volume 2. Elsevier Science, Amsterdam: North-Holland, 1989.

U. Doraszelski and A. Pakes. A Framework for Applied Dynamic Analysis in I.O. In M. Armstrong and R. Porter, editors, *Handbook of Industrial Organization*, volume 3. Elsevier, Amsterdam: North-Holland, 2007.

P. B. Ellickson and S. Misra. Supermarket Pricing Strategies. *Marketing Science*, 27(5):811–828, 2008.

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

R. Gilbert. Looking for Mr Schumpeter: Where Are We in the Competition-Innovation Debate? In A. B. Jaffe, J. Lerner, and S. Stern, editors, *Innovation Policy and the Economy (NBER)*, volume 6. The MIT Press, Cambridge: Massachusetts, 2006.

R. Goettler and B. Gordon. Competition and innovation in the microprocessor industry: Does AMD spur Intel to innovate more? Working Paper, 2009.

P. K. Goldberg. Product Differentiation and Oligopoly in International Markets: The Case of the U.S. Automobile Industry. *Econometrica*, 63(4):891–951, 1995.

P. K. Goldberg and F. Verboven. The Evolution of Price Dispersion in the European Car Market. *Review of Economic Studies*, 68(4):811–48, 2001.

Z. Griliches. Patent Statistics as Economic Indicators: A Survey. *Journal of Economic Literature*, 28(4):1661–1707, 1990.

B. H. Hall. Exploring the patent explosion. *Journal of Technology Transfer*, 30(1–2):35–48, 2004.

M. Jenkins, P. Liu, R. Matzkin, and D. McFadden. The Browser War. Econometric Analysis of Markov Perfect Equilibrium in Markets with Network Effects. U.C. Berkeley WP, 2004.

S. Klepper. The Capabilities of New Firms and the Evolution of the US Automobile Industry. *Industrial and Corporate Change*, 11(4):645–666, 2002.

T. Kretschmer, E. Miravete, and J. Pernias. Competitive Pressure and Innovation Complementarities. Working Paper, 2009.

R. C. Levin, W. M. Cohen, and D. C. Mowery. R&D Appropriability, Opportunity, and Market Structure: New Evidence on Some Schumpeterian Hypotheses. *American Economic Review*, 75: 20–24, 1985.

A. Pakes, M. Ostrovsky, and S. Berry. Simple Estimators for the Parameters of Discrete Dynamic Games, with Entry/Exit Examples. *Rand Journal of Economics*, 38(2), 2007.

M. Pesendorfer and P. Schmidt-Dengler. Asymptotic Least Squares Estimators for Dynamic Games. *Review of Economic Studies*, 75(3):901–928, 2008.

A. Petrin. Quantifying the Benefits of New Products: The Case of Minivan. *Journal of Political Economy*, 110(4):705–29, 2002.

S. Ryan. The Cost of Environmental Regulation in a Concentrated Industry. Working paper, 2006.

F. M. Scherer. Market Structure and the Employment of Scientists and Engineers. *American Economic Review*, 57(3):524–531, 1967.

J. A. Schumpeter. *Capitalism, Socialism and Democracy*. Harper, New York, 1942.

R. Siebert and C. Zulehner. The impact of market demand and innovation on market structure. Working paper, 2008.

J. Sutton. Market Structure: Theory and Evidence. In M. Armstrong and R. Porter, editors, *Handbook of Industrial Organization*, volume 3. North Holland, Amsterdam, 2008.

J. Van Biesebroeck. Productivity dynamics with technology choice: An application to automobile assembly. *Review of Economic Studies*, 70(1):167–198, January 2003.

X. Vives. Innovation and Competitive Pressure. *Journal of Industrial Economics*, 56(3):419–469, 2008.

G. Y. Weintraub, C. L. Benkard, and B. V. Roy. Markov Perfect Industry Dynamics with Many Firms. *Econometrica*, 76(6):1375–1411, 2008.

D. Y. Xu. A structural empirical model of R&D, firm heterogeneity, and industry evolution. Working Paper, 2008.

Figure 1: Aggregate relationship between innovation and market structure
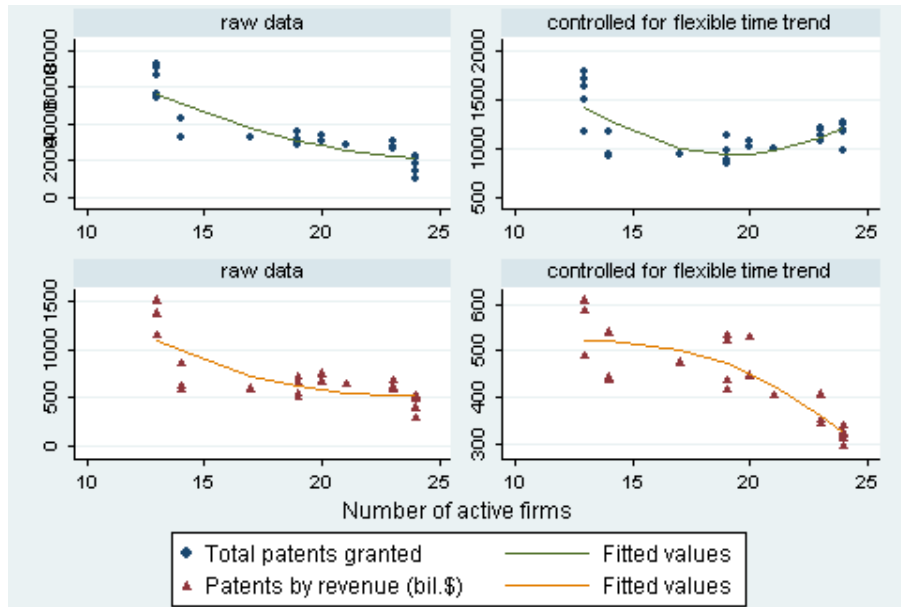

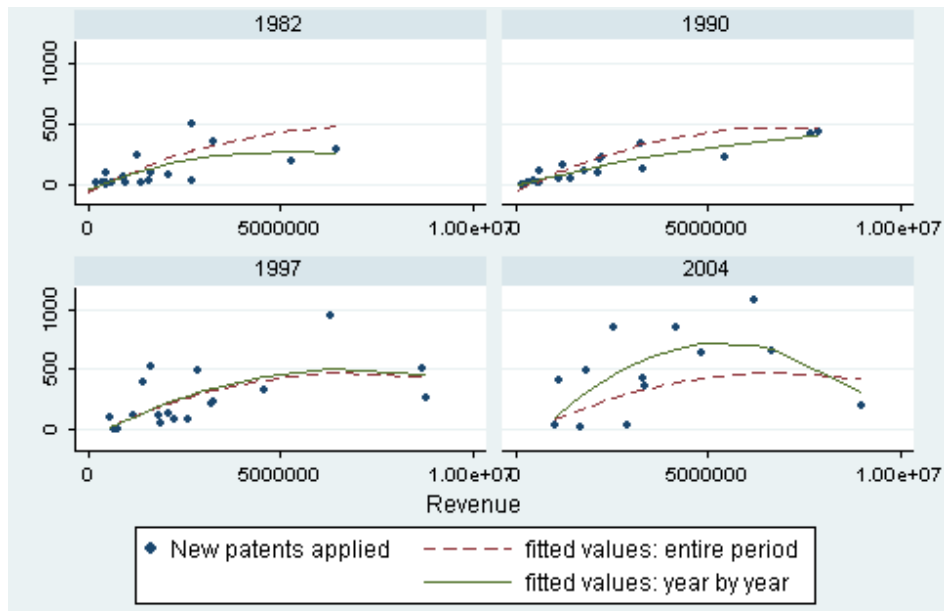Figure 2: Firm-level relationship between innovation and competitive position

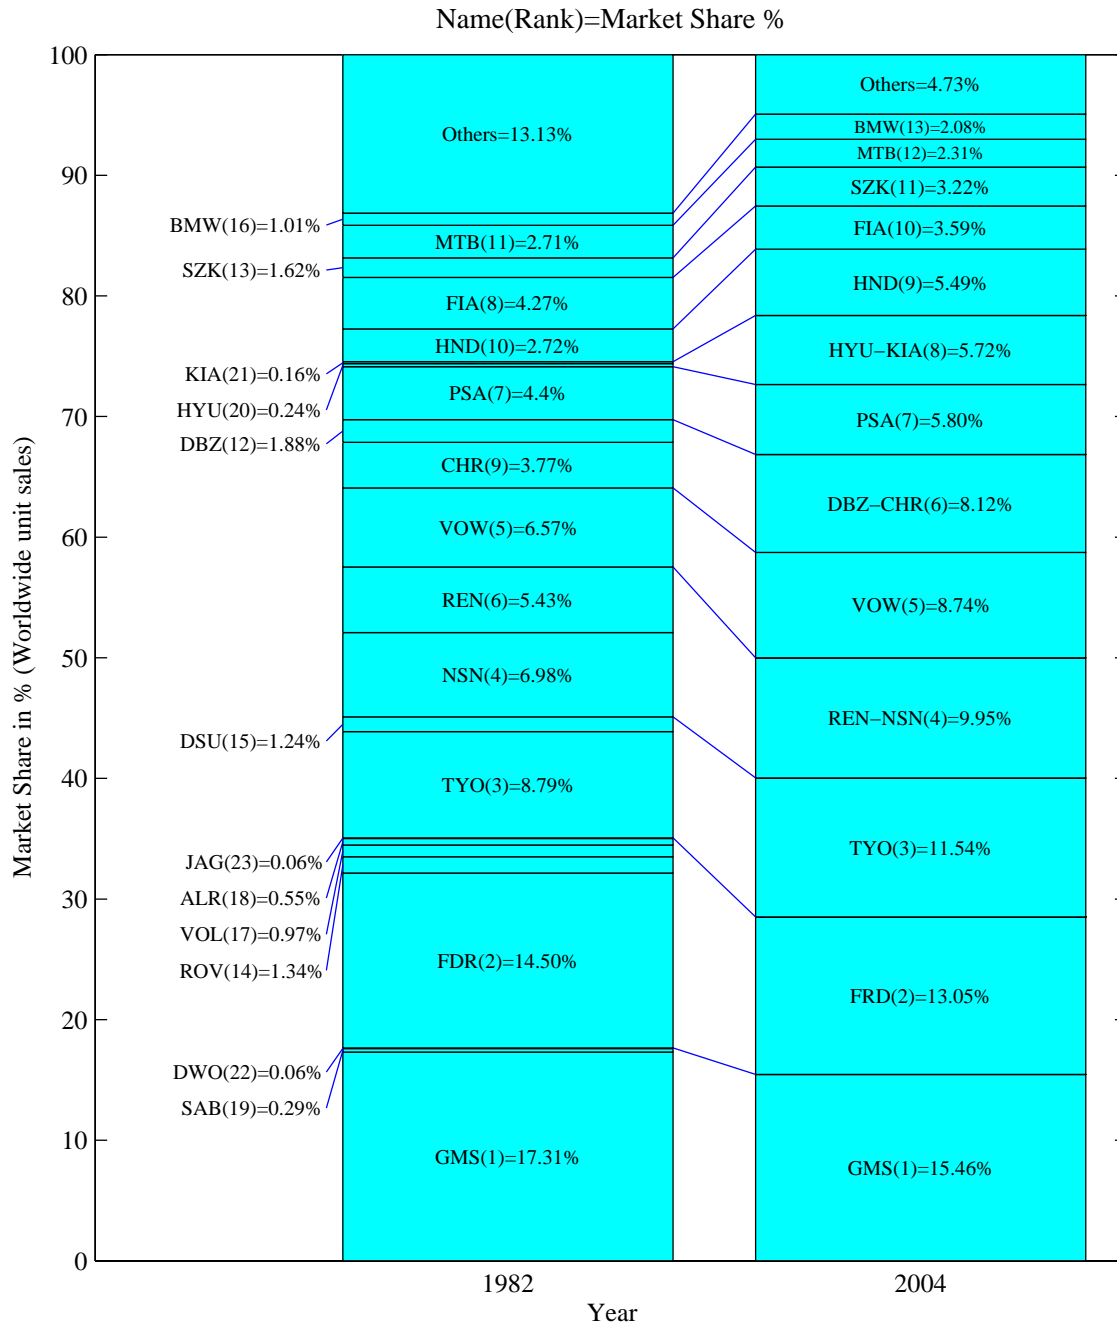Figure 3: Market shares of the firms in the sample

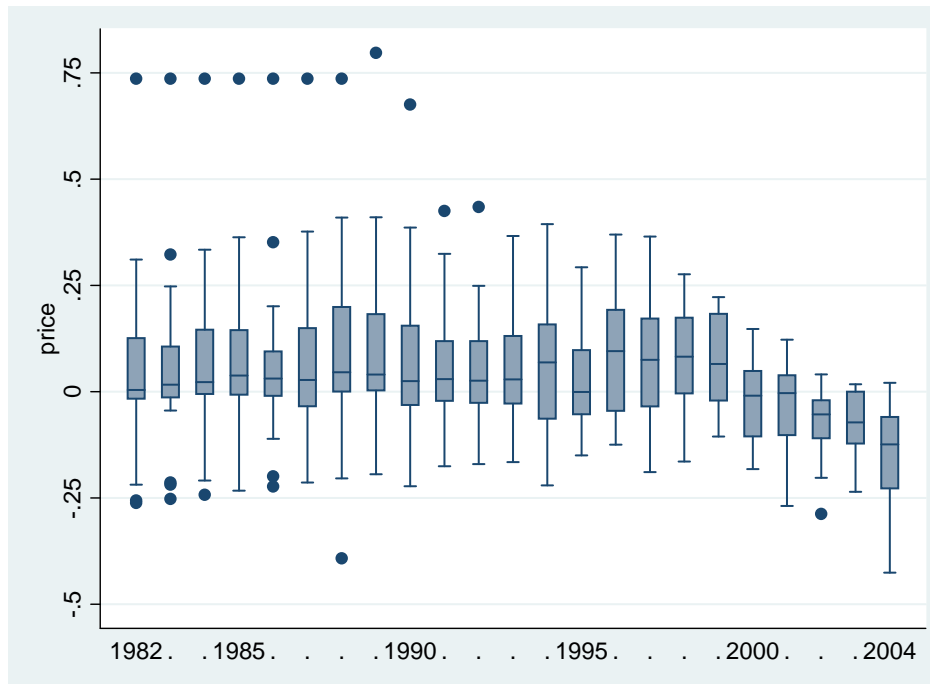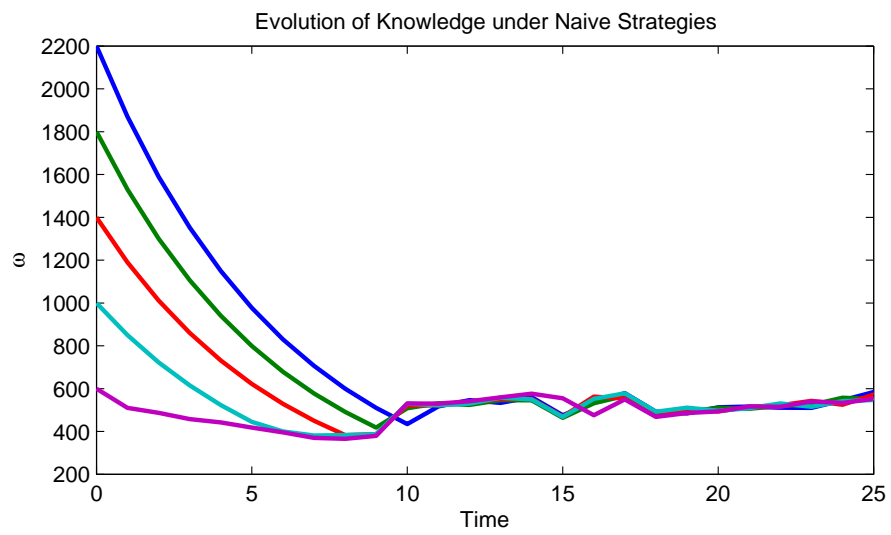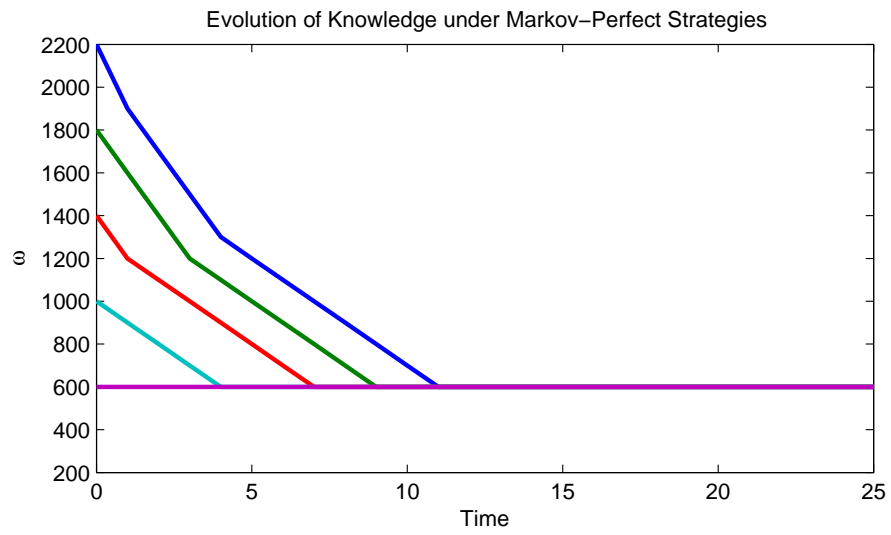Figure 4: Distribution of the hedonic prices relative to GM

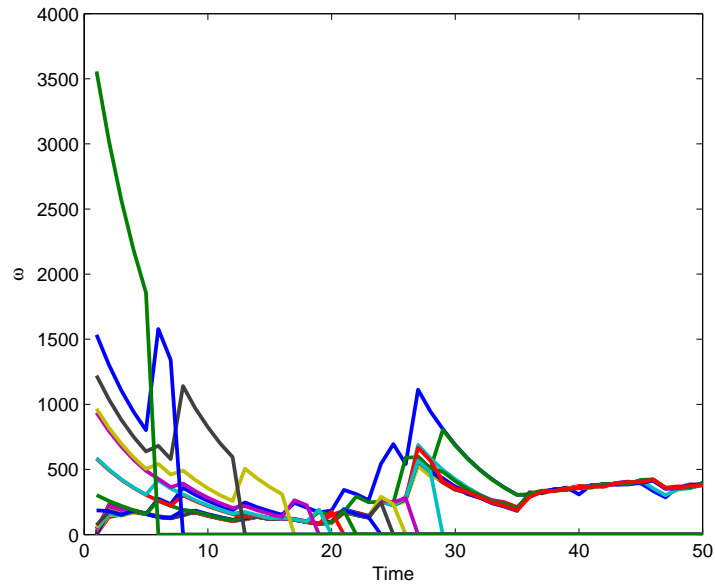Figure 5: A comparison of industry dynamics under MPE and naive strategies

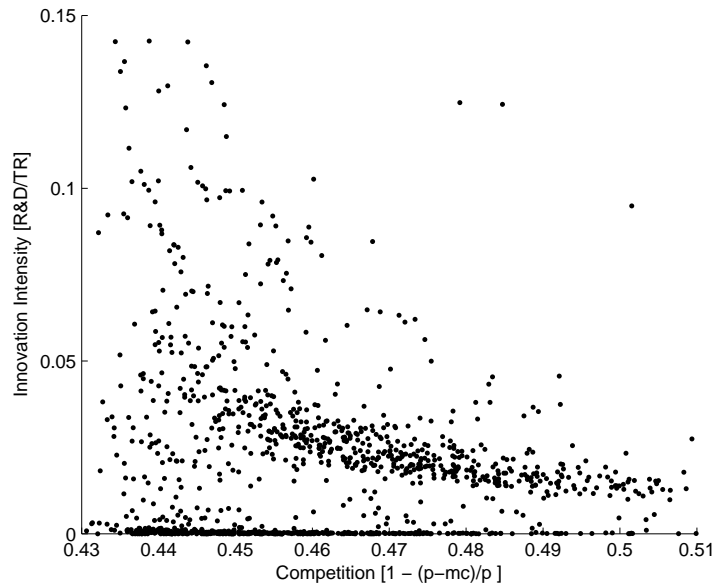Figure 6: Equilibrium Dynamics: Evolution of knowledge stock (by firm)

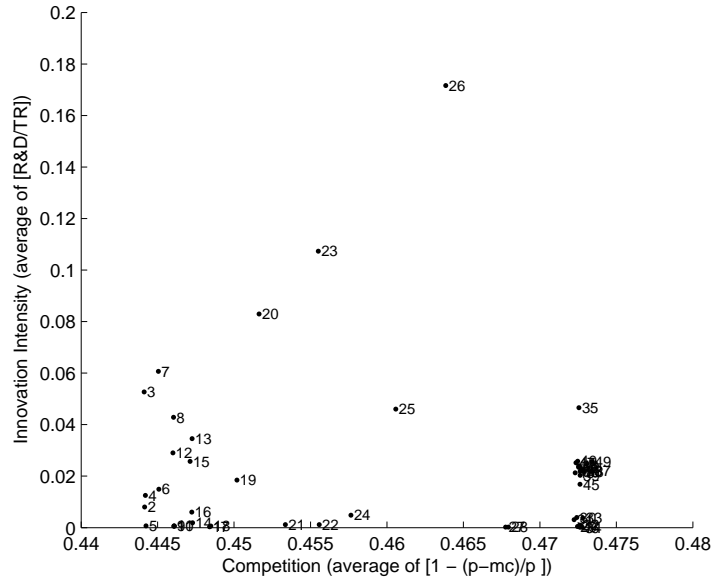Figure 7: Equilibrium Dynamics: Relationship between Competition and Innovation at the Firm Level

Figure 8: Equilibrium Dynamics: Relationship between Competition and Innovation at the Industry Level
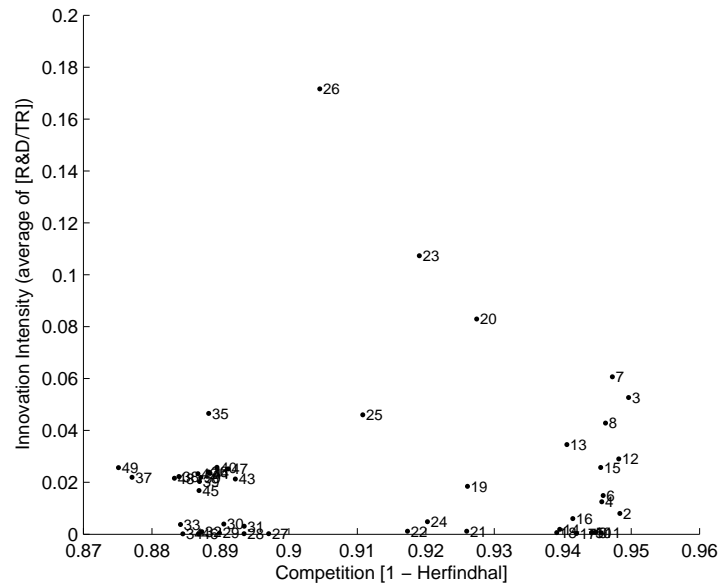

Figure 9: Equilibrium Dynamics: Relationship between Competition (1 - Herfindhal) and Innovation at the Industry Level

Table 1: R&D expenditure by industry in the most innovative countries

(data for 2006, in bil. PPP)

Industry (ISIC Rev. 3)	USA	Japan	Germany	Korea	France
Chemicals (24)	46.3	16.4	8.2	2.1	5.0
Radio, TV, telecommunications eq. (32)	31.2	12.2	4.1	13.3	2.8
Motor vehicles (34)	**16.6**	**17.9**	**14.4**	**4.2**	**4.6**
Medical, precision, optical instruments (33)	22.4	4.6	3.5	0.4	1.6
Office, accounting, computing mach. (30)	7.4	14.1	0.6	0.4	0.2

Note: All sectors which are in the top three by total R&D expenditure in any of the five countries

are included; industries are sorted by total R&D expenditure across the five countries.

Source: OECD ANBERD database, edition 2009 (online).

Table 2: Benchmark estimates for demand and marginal cost parameters

	Demand	Marginal cost
Price	-2.355	
	(0.337)	
Knowledge (ω)	0.166	0.427
	(0.046)	(0.087)
Knowledge-squared		-0.180
		(0.034)
Unobserved quality ($\tilde{\xi}$)		-0.115
		(0.049)
Unobserved quality-squared		0.006
		(0.004)

Joint estimation of demand and first order condition, with firm fixed

effects and AR(1) process in the demand equation and using

competitors' knowledge as instrument for price.

Table 3: Sensitivity analysis for demand and marginal cost parameters

| | Demand parameters | | | |
	(1)	(2)	(3)	(4)
Price	-1.742	-2.175	-2.421	-2.534
	(0.083)	(0.284)	(0.357)	(0.375)
Knowledge (ω)	0.195	0.385	0.187	0.168
	(0.036)	(0.023)	(0.047)	(0.048)
	Marginal cost parameters			
	(1)	(2)	(3)	(4)
Knowledge (ω)	0.017	-0.086	0.083	0.058
[in logs for (3)]	(0.092)	(0.099)	(0.015)	(0.024)
Knowledge-squared	-0.052	0.020		
	(0.040)	(0.038)		
Unobserved quality ($\tilde{\xi}$)	-0.155	-0.439	-0.268	-0.024
[in logs for (3)]	(0.053)	(0.148)	(0.042)	(0.013)
Unobserved quality-squared	0.015	0.035		
	(0.007)	(0.004)		

All columns contain results of joint estimation of demand and first order condition.

(1) Estimation is with OLS, i.e. without instruments for price.

(2) Estimation is with IV, but unobserved quality ξ is assumed to be constant.

(3) Estimation as in Table 2, but MC is a function of the log of the state variables.

(4) Estimation as in Table 2, but MC is a linear function of the state variables.

Table 4: Sensitivity analysis for the dynamic parameter

(a) Varying demand		(b) Varying MC	
IV	21.8	quadratic	21.8
OLS	31.2	linear	22.7
constant ξ	28.3	log-linear	26.3

(c) Varying policy		(d) Varying ξ_{A+B}	
non-parametric	21.8	average A & B	21.8
restricted (8 terms)	14.3	maximum A or B	28.8
non-parametric in logs	21.0	ω-weighted average	27.0

(e) Discount factor		(f) Depreciation rate	
benchmark (0.92)	21.8	benchmark (0.15)	21.8
low (0.90)	20.7	low (0.10)	18.1
high (0.94)	23.8	high (0.20)	25.2

(g) EU-US patent ratio	
benchmark (2.36)	21.8
low (1.18)	26.8
high (3.54)	13.0

A   Appendix: The Merger Technology

This Appendix accomplishes two things. First, we write down the value function of a firm when it has more than one rival. Second, we calibrate the probability of a merger in such a way that if the model is simulated forward, the simulated number of mergers would, on the average, equal the rate at which we observe mergers over the sample period.

With more than two firms in the industry and the possibility of mergers, the value function of firm j is a generalization of (9). For expositional convenience, we assume below that there cannot be more than one merger in a period. Let $p_m(n)$ denote the probability that a firm experiences a merger at the end of the current period. This probability is an increasing function of the number of firms in the industry (n) and constant across firms. k indexes a firm that j is randomly matched with to merge and $\mathbf{s}'^{-jk}$ is the industry state excluding firms j and k after the merger. The value function of firm j can then be written as

$$
\begin{aligned}
V_j(\omega_j, \xi_j, \mathbf{s}^{-j}) &= \max_{x \in \mathbb{R}^+} \Bigg\{ \pi_j(\omega_j, \xi_j, \mathbf{s}^{-j}) - cx(\omega_j, \xi_j, \mathbf{s}^{-j}) \\
&\quad + \beta \Big[\frac{p_m}{n-1} \sum_{k \neq j} \zeta_j(\cdot) E[V_j(\omega_j' + \omega_k', (\xi_j' + \xi_k')/2, \mathbf{s}'^{-jk})] \\
&\quad + \frac{(1-p_m)\left[1 - (1-p_m)^{n-1}\right]}{(n-1)(n-2)/2} \sum_{k \neq j} \sum_{l \neq k,j} E[V_j(\omega_j', \xi_j', \omega_k' + \omega_l', (\xi_k' + \xi_l')/2, \mathbf{s}'^{-jkl})] \\
&\quad + (1-p_m)^n E[V_j(\omega_j', \xi_j', \mathbf{s}'^{-j})] \Big] \Bigg\},
\end{aligned}
\tag{17}
$$

The expression for the continuation value contains three terms. The probability that firm j will be involved in a merger is p_m. Since this merger could be with any of its $n-1$ rivals, the probability of merging with any particular rival is $\frac{p_m}{n-1}$. The first term sums over the $n-1$ merger possibilities, while updating the state variables for the affected firms. The fraction of the total value of the merged firm that is accrues to V_j is ζ_j, which is still given by (10).

The probability that firm j will not be involved in a merger the next period is $1 - p_m$. Since it is the same for all n firms, the probability that there will be no merger in the industry is $(1 - p_m)^n$. The third term captures the next period valuation for firm j in this case.

The probability that there is a merger in the industry, but firm j is not a part of it must thus be

$$
1 - [p_m + (1-p_m)^n] = (1 - p_m)\left[1 - (1-p_m)^{n-1}\right]. \quad (*)
$$

How many such mergers are possible? The total number of possible mergers (regardless of whether firm j is in a merger or not) is $\frac{n^2 - n}{2}$. This is the number of elements in the lower-diagonal submatrix of an $n \times n$ matrix. Out of these, firm j could be part of $n - 1$ mergers. Hence the

42

number of possible mergers that do not involve firm j is

$$\frac{n^2 - n}{2} - (n-1) = \frac{(n-1)(n-2)}{2}. \quad (**)$$

Dividing $(*)$ by $(**)$ gives the individual probability of each of those mergers, not involving firm j. The second term in the expression for the continuation value of firm j sums the valuation for firm j in the case of each possible merger of two rivals k and l, weighted by the (constant) probability of each individual merger.

As discussed earlier, we set a lower bound on the number of firms in the industry at ten. At the lower bound $p_m = 0$ and (17) reduces to (6).

We now show how we can impute the probability that a firm will be involved in a merger (p_m), from the observed number of mergers in the data. Let p be the probability (same for all firms) that a firm will be up for merger this period, which differs from the merger probability p_m, as mergers only take place if at least two firms are up for merger.

With only two active firm, $p_m = p^2$. With more than two active firms, the merger probability is the sum of the following probabilities (if more than two firms are up for merger, we pick two at random):

$$
\begin{aligned}
\Pr(\text{The firm is up for merger and 1 other firm is also up}) &= pP_1^{n-1} \\
\Pr(\text{The firm is up for merger and 2 other firms are up}) &= \frac{2}{3}pP_2^{n-1} \\
\Pr(\text{The firm is up for merger and 3 other firm are up}) &= pP_3^{n-1} \\
\Pr(\text{The firm is up for merger and 4 other firms are up}) &= \frac{4}{5}pP_4^{n-1} \\
\vdots &= \vdots
\end{aligned}
$$

where P_k^n is the binomial probability that out of n firms exactly k are up for mergers and is given by $\binom{n}{k}p^n(1-p)^{n-k}$. If n is odd, the last term in the above sequence would be $\frac{n-1}{n}P_{n-1}^{n-1}$; if n is even, the last term would be P_{n-1}^{n-1}. Adding all terms together, we can write the sum as

$$p_m(n,p) = \left[\sum_{i \in O, i < n} P_i^{n-1} + \sum_{i \in E, i < n} \frac{i}{i+1} P_i^{n-1} \right] \cdot p, \quad (18)$$

where $O = \{1, 3, 5, \ldots\}$ and $E = \{2, 4, 6, \ldots\}$.

We have derived an expression for p_m in terms of n and p. Next, we impute p from the data. In our sample we observe 10 mergers in 23 years, so the expected number of mergers in any period is $10/23$. If $\bar{n}$ is the average number of active firms, we can force the expected number of mergers to be a declining function of n, by defining it as:

$$E(M) = \frac{10}{23} \cdot \frac{n}{\bar{n}}. \quad (19)$$

43

The expected number of mergers is implicitly a function of p as follows:

$$E(M) = \sum_{M=1}^{\lfloor \frac{n}{2} \rfloor} M \cdot P(M),$$

where $P(M)$ is the probability that M mergers take place, e.g. $P(1) = (P_2^n + P_3^n)$, the probability that one merger will take place is simply the sum of the probabilities that out of n firms 2 or 3 are up for merger. This can explicitly be written as

$$E(M) = \begin{cases} \sum_{i=1}^{\lfloor \frac{n}{2} \rfloor - 1} i(P_{2i}^n + P_{2i+1}^n) + \lfloor \frac{n}{2} \rfloor (P_{2\lfloor \frac{n}{2} \rfloor}^n + P_n^n) & \text{if } n \text{ is odd} \\ \sum_{i=1}^{\lfloor \frac{n}{2} \rfloor - 1} i(P_{2i}^n + P_{2i+1}^n) + \frac{n}{2} P_n^n & \text{if } n \text{ is even.} \end{cases} \tag{20}$$

Equating (19) and (20) allows us to solve for p for each n. Once we have p, we can use (18) to calculate p_m.

B   Appendix: Additional estimates

Table 5: Parameters on the polynomials $(\omega_j)^k(\omega_{-j})^l(\xi_j)^m$ in the estimated policy function

α_{klm}	Dependent variable: number of new patent applications					
(klm)	(1)		(2)		(3)	
000	49.443	(0.49)	-180.2	(2.94)	718166	(4.56)
100	-0.242	(3.02)	0.137	(6.53)	-6469.16	(1.42)
010	-0.002	(0.22)	0.021	(3.45)	-215860	(4.51)
001	320.0	(0.93)	-14.92	(0.20)	2873.89	(0.29)
200	2.45e-5	(1.06)	-8.57e-6	(1.79)	-233.087	(2.75)
020	-1.28e-7	(0.51)	-5.55e-7	(3.77)	21757.8	(4.47)
002	491.6	(1.90)	-7.780	(0.10)	1766.44	(3.58)
110	4.38e-5	(6.84)	3.84e-6	(3.87)	1284.05	(1.39)
101	-0.251	(1.52)	0.025	(0.85)	-1095.49	(2.83)
011	-0.034	(1.17)			700.370	(0.35)
021	1.52e-6	(2.33)			-100.754	(0.94)
012	-0.041	(2.77)			-170.704	(3.25)
120	-0.83e-9	(5.69)			-62.115	(1.30)
210	-1.57e-9	(2.07)			18.837	(1.87)
102	0.155	(1.71)			-10.207	(0.53)
201	1.31e-5	(0.50)			0.387	(0.05)
111	1.69e-6	(0.22)			110.564	(2.55)
300					5.322	(3.70)
030					-735.848	(4.43)
003					12.727	(0.67)

Estimation by OLS; absolute value of t-statistics are in parenthesis. In columns (1) and (2), the variables in the polynomials are in levels; in column (3) they are in logarithms.