Market Structure and Innovation:
A Dynamic Analysis of the Global Automobile Industry*

Aamir Rafique Hashmi†
and
Johannes Van Biesebroeck‡

February 12, 2007

Abstract

The question that how market structure and innovation are related has been extensively studied in the literature. However, there is hardly any notable study on this question for the global automobile industry. We fill this gap by studying the relationship between market structure and innovation in the global automobile industry for the 1980-2005 period. We use the dynamic industry framework of Ericson and Pakes [1995] and estimate the parameters of the model using a two-step procedure proposed by Bajari et al. [2006].

Since the global auto industry has seen a lot of consolidation since 1980, mergers are an important ingredient of our model. After estimating the parameters of the model, we simulate the industry forward and study how changing market structure (mainly due to mergers) affects innovative activity at the firm as well as the industry level. Our findings are the following. (1) The effect of market structure on innovation in the global auto industry depends on the initial state of the industry. If the industry is not very concentrated, as it was in 1980, some consolidation may increase the innovative activity. However, if the industry is already concentrated, as in 2005, further consolidation may reduce the innovation incentives. (2) Mergers reduce the value of merging firms though they may increase the aggregate value of the industry. (3) Mergers between big firms eventually reduce consumers' utility.

Key words: Competition and Innovation; Automobile Industry; Dynamic Games
JEL Classification Codes: C73; L13; L62; O31

---

*We thank participants at the 2006 annual meetings of the Canadian Economics Association (Montreal) and the European Economic Association (Vienna) and the CEPA seminar at the University of Toronto for their comments. Financial support from AUTO21, CFI, and SSHRC is greatly appreciated.

†Departments of Economics, National University of Singapore and University of Toronto. E-mail: aamir.hashmi@gmail.com

‡Department of Economics, University of Toronto and NBER. E-mail: johannes.vanbiesebroeck@utoronto.ca

1 Introduction

In this paper, we investigate how changes in industry structure interact with the innovative activity of the incumbent firms in the global automobile industry. The question of interaction between market structure and innovation has been the subject of intense research since Schumpeter [1950].[1] However, there is hardly any study on the global automobile industry. This is surprising given the fact that the global automobile industry has not only undergone interesting changes in its market structure in recent years but is also one of the most innovative industries in the global economy.

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. On the one hand, this consolidation is the result of increased competition, which has made it harder for smaller firms to survive on their own. On the other hand, consolidation intensifies competition as the emerging groups are highly research intensive. The top 13 firms in the auto industry spent more than 55 billion dollars on R&D in 2005 and have obtained more than 50,000 patents from the US patent office between 1980 and 2004.

Automobile industry is a highly research intensive industry: in 2003, 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 illustrate the comparative importance of R&D in the automotive industry for the OECD and the three largest economic blocs. The industry is concentrated worldwide, making it highly likely that the main firms will take actions of competitors into account when deciding on their own innovative activities. In 2005, more than 85% of all vehicles were produced by the 13 largest firms, which are active in all major regions of the world.[2]

[Table 1 approximately here]

The reasons for auto industry being so innovation intensive are simple. At the most fundamental level, two features are necessary for innovation to take place. First, firms need to be able to finance innovation, i.e. there needs to be a margin between price and marginal cost. Second, firms need to have an incentive to innovate, i.e. innovation has to increase expected profits. Both conditions are clearly satisfied in the automobile industry. Demand estimates for this industry — see for example Berry et al. [1995] or Goldberg [1995] — reveal that markups over marginal costs tend

[1] Kamien and Schwartz [1982], Cohen and Levin [1989], Ahn [2002] and Gilbert [2006] are surveys of the relevant literature.

[2] Throughout, we do not distinguish between minority share holdings and outright control. E.g. Mazda is counted as part of the Ford group, even though Ford Motor Co. never held more than 33.4% of Mazda's shares (achieved by 1996).

to be large, consistent with the view that fixed costs are important in this industry. Innovation is also likely to advance a firm's competitive position through higher product quality and reliability, the introduction of desirable features in its products, or lower production cost. A large number of analysts makes a living measuring these benefits for consumers and investors alike.[3]

Unlike most of the studies that use reduced form regressions to study the relationship between market structure and innovation, we do so in a strategic and dynamic environment. There are at least two reasons for this. First, innovation is inherently a dynamic activity. Firms make R&D investments today expecting uncertain future rewards through better products or more efficient production. The magnitude of these benefits are intricately linked to the future market structure and the level of innovation of competitors.[4] Merging or forming strategic alliances with rivals is also best analyzed in a dynamic framework. Mergers interact with innovation through their influence on market structure as well as through the consolidation of the knowledge stock in the industry. The global automobile industry seems an ideal place to study these two forces and their interaction.

Second, we now have methods to estimate models of dynamic competition. In this study, we employ a recently developed technique in the estimation of dynamic games that does not require one to solve for the equilibrium of the game, see Bajari et al. [2006]. Our study is one of the first to put these new techniques to a practical test, demonstrate their usefulness and highlight some practical difficulties associated with their application.[5]

The principal objective of our paper 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 so in a dynamic environment in which firms produce differentiated products that differ in quality. Firms invest in R&D to improve the quality of their products and to lower the cost of production. The market share of a firm depends on the relative quality of its product and the price, which is set strategically in each period. Investments in R&D increase the technological knowledge of the firm but the exact outcome of R&D is uncertain. On average, higher knowledge translates into a higher quality product and lower marginal cost. The investment in R&D is modeled as a strategic decision: a firm takes the actions of its rivals and their possible future states into account before making its R&D decision. In addition, firms incorporate potential future mergers into their expectations.

[3]To name but a few, *J.D. Power* and *Consumer Reports* measure defect rates in vehicles; numerous consumer magazines and internet sites compare the performance and discuss the features of vehicles; and *Harbour Consulting* and *KPMG* continuously compare productivity in the industry.

[4]Aghion and Griffith [2005] provide an overview of the issues involved and review some popular modeling approaches.

[5]Among a growing list of (working) papers using various approaches, we can point the interested reader to the following applications: Ryan [2006], studying the effect of environmental regulation on the cement industry; Collard-Wexler [2005], studying the effect of demand fluctuation in the ready-mix concrete industry; Aguirregabiria and Ho [2006], studying the airline industry; and Sweeting [2006] estimating switching costs for radio station formats.

First, we obtain an estimate of our only dynamic parameter – the cost of innovation – which can only be pinned down in a dynamic game-theoretic model. Preliminary estimates put the R&D cost to obtain one new patent (in expectation) at \$29.9m if marginal costs decline linearly with knowledge or at \$11.6m if we limit the model to product innovations, i.e. marginal costs are constant. These numbers seem reasonable, given that the observed median for the R&D to patent ratio in our sample is \$14.9m and the mean is \$15.6m.

Second, with parameter estimates in hand, we study how changes in market structure affect the intensity of innovation. Our main finding is that if the industry is fragmented, some consolidation may improve the innovation intensity. However, once the industry is concentrated enough, any further consolidation will only reduce innovation incentives.

Third, we study the effects of mergers on firm values and on consumers' utility. We find that mergers destroy some value of the merging firms. However, they help rivals by reducing the extent of competition in the market. The overall effect of mergers on aggregate industry value is positive if merging firms are large in size. These large mergers, however, have negative effects on consumers' utility. On the other hand, consumers seem to benefit when relatively smaller firms merge.

The remainder of the paper is organized as follows. The supply and demand side of the model as well as the Markov perfect equilibrium concept we rely upon are introduced in Section 2. Section 3 introduces the data. The two-step estimation methodology is discussed in Section 4: In the first step, we estimate all the static parameters and generate the value functions using forward simulation. In the second step, we estimate the only dynamic parameter of the model. Section 5 discusses the estimation results. The impact of changes in market structure on incentives for innovation, firm value and consumer utility is in Section 6 and Section 7 concludes.

2 The Model

Our modeling strategy follows Ericson and Pakes [1995]. There are n firms, each producing a differentiated vehicle. Firms differ in their technological knowledge, which is observable by all market participants as well as by the econometrician, and in some product-specific characteristics, which are known to the market participants but are not observed by the econometrician. We denote the technological knowledge of firm j ($j = 1, 2, \ldots, n$) by $\omega_j \in \mathbb{R}^+$ and the index of product-specific characteristics of firm j by $\xi_j \in \mathbb{R}$. For the industry as a whole, the vectors containing ω and ξ are denoted by $\mathbf{s}_\omega$ and $\mathbf{s}_\xi$.[6] For later use, we define the vectors of ω and ξ excluding the firm j as ${\mathbf{s}_\omega}^{-j}$ and ${\mathbf{s}_\xi}^{-j}$.[7] We also define $\mathbf{s} = \{\mathbf{s}_\omega, \mathbf{s}_\xi\}$ and $\mathbf{s}^{-j} = \{{\mathbf{s}_\omega}^{-j}, {\mathbf{s}_\xi}^{-j}\}$.

Time is discrete. At the beginning of each period, firms observe $\mathbf{s}$ and make their pricing and

[6] $\mathbf{s}_\omega = [\omega_1\ \omega_2\ \ldots\ \omega_n]$ and $\mathbf{s}_\xi = [\xi_1\ \xi_2\ \ldots\ \xi_n]$.

[7] ${\mathbf{s}_\omega}^{-j} = [\omega_1\ \omega_2\ \ldots\ \omega_{j-1}\ \omega_{j+1}\ \ldots\ \omega_n]$ and ${\mathbf{s}_\xi}^{-j} = [\xi_1\ \xi_2\ \ldots\ \xi_{j-1}\ \xi_{j+1}\ \ldots\ \xi_n]$.

investment decisions (see below). Although the pricing decisions are static, the investment decisions are dynamic and depend on the current as well as expected future states of the industry.

Firms make investment to increase their technological knowledge. A higher level of knowledge may boost demand, for example, by improving vehicle quality or by introducing more innovative product features. It may also reduce marginal cost by improving the efficiency of production. Hence, the model features both product and process innovation. The effect of R&D investment on knowledge is the sum of a deterministic and a random component, capturing that innovation is a stochastic process. The technological knowledge depreciates at an exogenous rate.

There is no entry or exit in our model, reflective of the evolution of the industry over the last 25 years. However, a firm may 'exit' by merging with another firm. Mergers are an important component of our model as they lead to discrete changes in the market structure and force the firms to readjust their prices and investments to take new industry structure into account. 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.

2.1 The Demand Side

Following Berry et al. [1995] and several others studying the automobile industry, we use a discrete choice model of individual consumer behavior to model the demand side.[8] There are n firms in the industry producing differentiated vehicles. Vehicles differ in quality, which has two components. The first component, observable to the market participants as well as the econometrician, is positively related to the firm's technological knowledge, $g(\omega)$, with $(\partial g(\omega)/\partial \omega \geq 0)$. The second component ($\xi$) is unobservable to the econometrician but firms and consumers know it and use it in their demand and pricing decisions.

There are m consumers in the market in each period and each of them buys one vehicle. m grows at a constant rate that is exogenously given. The utility of a consumer depends on the quality of a vehicle, its price and the consumer's idiosyncratic preferences. The utility consumer i gets from buying vehicle j is

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

$g(\omega_j)$ is the observable quality of the vehicle produced by firm j; for simplicity, we assume that it equals $\log(\omega_j + 1)$. p_j is the price of the vehicle produced by firm j, ξ_j is the unobservable quality and ν_{ij} is the idiosyncratic utility that consumer i gets from good j. θ_ω and θ_p are preference parameters. θ_ω shows how quality conscious the consumers are and θ_p is a measure of their price elasticity. We assume ν are i.i.d. extreme value distributed, which gives the following expected

[8] Our description of the multinomial logit model follows Anderson et al. [1992], pp. 39-40.

market share for firm j:

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

where p_j is the price charged by firm j and $\mathbf{p}^{-j}$ is the price vector of all the other firms (excluding firm j) in the industry. The expected demand for vehicle j 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 because its equilibrium price is a function of its rivals' prices.

2.2 The Supply Side

We begin our explanation of the supply side with the static profit function. We assume that R&D investments will only generate useful knowledge with a one period lag and that prices can be adjusted flexibly period by period. Hence, at the beginning of each period, after observing their individual and industry states, the firms engage in a differentiated products Bertrand-Nash game. Each firm chooses its own price to maximize profits, taking the prices of its rivals ($\mathbf{p}^{-j}$) and industry state ($\mathbf{s}$) as given.

The profit maximization problem of an individual firm j is

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

where mc_j is the marginal cost incurred by firm j to produce a vehicle. The marginal cost is a function of the firm's knowledge, capturing cost reducing process innovations. fc_j is the fixed cost of operations faced by firm j. For now, the fixed cost does not play any role and we simply set it equal to zero.

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

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

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

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

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

[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 R&D investment based on the expected value of future profit streams. The problem is recursive and can be described by the following Bellman equation

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

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

The knowledge of the firm evolves according to the following equation:

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

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

We assume an exogenous AR(1) process for the evolution of unobserved component of quality (ξ).

$$\xi'_j = \rho \xi_j + \epsilon_{\xi_j}. \tag{8}$$

We now add mergers to the model. The first question in this regard is about merger technology: why mergers take place and what are the characteristics of the firms that merge? During the 26 year period we study (1980–2005), a total of 9 mergers, acquisitions and associations took place among the firms in our sample. These were: Ford-Jaguar (1989); GM-Saab (1990); Ford-Volvo (1998); DaimlerBenz-Chrysler (1998); Hyundai-Kia (1998); Toyota-Daihatsu (1999); Renault-Nissan (1999); and Ford-Rover (2000). Table 2 indicates how these events have reshaped the structure of the industry.[11]

There is no clear pattern in the merger activity, making it nearly impossible to predict. 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.[12] In several instances, a larger firm has taken over a smaller one, which is not unexpected in an industry with large scale economies (certainly at the firm level). Examples are Ford's acquisition of Jaguar, Land Rover, and Volvo;

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

[11] Prior to our sample, in 1979, Ford and Mazda entered into an 'international partnership', with Ford taking a 25% stake (later increased to 33.4% in 1996) in Mazda.

[12] This happened in spite of a Ford ownership interest (since 1986).

GM's takeover of Saab and Daewoo; and Toyota's takeover of Daihatsu. Such attempts are not always successful. Several alliances were tried out before a complete merger was attempted, but fell apart, e.g. DaimlerChrysler with Mitsubishi; GM with Suzuki and Isuzu; and Honda with Rover. There are also several smaller independent companies that are thriving in spite of incessant merger predictions. Honda, BMW, and Porsche are consistently among the most profitable firms in the industry.

There have also been instances of two large firms merging or 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. Even among larger firms, attempts to cooperate sometimes fail as with Renault and AMC (1979–1987) and with DaimlerChrysler and Hyundai (2000–2004).

It is especially difficult to model mergers is a simple environment like ours in which there are only two state variables. In the model, the benefits of a higher level of knowledge —product and process innovations that boost demand and reduce costs—can be captured instantaneously. This precludes as a motive for merging the combination of a large firm that has many models and dealerships with a smaller firm that has a higher level of technological knowledge. Given that the policy function is estimated concave in a firm's own knowledge (see Section 5.1.3), gaining scale economies in R&D or product development can also not explain mergers in our model.

In the absence of any clear pattern in mergers, it seems reasonable to assume that mergers are random and occur with an exogenously given probability. We pick this probability such that the number of mergers generated by the model are, on the average, equal to the actual number of mergers observed in the data. However, if we allow firms to continue to merge this way, in the limit the industry will approach a monopoly, which is highly unlikely in reality. By defining the expected number of mergers as a declining function of n (see the calibration of the merger probability in the Appendix), we have lengthened the time it takes for the industry to become a monopoly, but in the limit this does not solve the problem.

There are at least two solutions to the above problem. First, on competition policy grounds we may impose a floor on the number of firms in the industry. Once the number of firms in the industry hits the lower bound, no future mergers are allowed. This will be in line with Klepper [2002a] and Klepper [2002b], which argue that the U.S. automobile industry has settled into a stable oligopoly. The second solution would be to introduce entry, but this would require more radical changes in our empirical strategy. In practice, entry has hardly played any role in the evolution of the global market structure, although this might change with the development of the Chinese and Indian automotive industries. Hence, we adopt the first solution and set a floor on the number of firms in the industry.

Even with this simple merger technology we need to specify the values of the state variables for the merged entity and how each firm incorporates possible mergers in the evaluation of its future value. With regard to the first question, we assume that when two firms merge, the knowledge of the new firm is the sum of the individual knowledge stocks. Another possibility would be to allow for complementarities in the knowledge stock or, at the other extreme, assume some overlap in knowledge and discount the sum. All three assumptions are equally arbitrary, but we feel that simply adding the two knowledge stocks is closest in spirit to the state transition function for knowledge used throughout. As for the second state variable (ξ), we assume that the unobserved quality of the vehicle produced by the merged firm is the average of the unobserved qualities of the vehicles produced by the individual firms before the merger. This assumption finds some support in the data when we estimate ξ's as the residuals from our demand equation (see below).

To answer the second question, first consider the situation where A and B are the only firms in the industry and their respective states are (ω_A, ξ_A) and (ω_B, ξ_B). As long as they are independent firms, there is an exogenously given probability p_m each period that they will merge. How should firms A and B incorporate this information into their calculations of the value functions? The value function for firm A, the result for firm B follows by symmetry, becomes

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

where $\zeta_A(\omega'_A, \xi'_A, \omega'_B, \xi'_B)$ is the share of firm A in total value of the merged firm. We assume that this share is simply the ratio of the value of the firm to sum of the values of both firms in absence of the merger:

$$\zeta_A(\omega'_A, \xi'_A, \omega'_B, \xi'_B) = \frac{EV(\omega'_A, \xi'_A, \omega'_B, \xi'_B)}{EV(\omega'_A, \xi'_A, \omega'_B, \xi'_B) + EV(\omega'_B, \xi'_B, \omega'_A, \xi'_A)}.$$

The same idea extends to an industry with more than two firms, but the computations become more involved. Let $p_m(n)$ be the probability that a firm experiences a merger at the end of the current period. This probability is an increasing function of n and constant across firms.[13] Let j be the index of the firm that we are interested in and let k be the index of its rivals ($k \neq j$ i.e. $k = 1, 2, \ldots, j-1, j+1, \ldots, n-1$). Let $\mathbf{s}'^{-jk}$ be the industry state (excluding firms j and k) after

[13] When the number of firms in the industry declines, the probability of merger also declines.

the merger between firms j and k. The value function of firm j can then be written as

$$\begin{aligned} V_j(\omega_j, \xi_j, \mathbf{s}^{-j}) \quad &= \quad \max_{x_j \in \mathbb{R}^+} \Big\{ \pi_j(\omega_j, \xi_j, \mathbf{s}^{-j}) - c x_j(\omega_j, \xi_j, \mathbf{s}^{-j}) \\ &\quad + \beta \Big[\frac{p_m}{n-1} \sum_{k \neq j} \zeta_j(\cdot) E V_j(\omega_j' + \omega_k', (\xi_j' + \xi_k')/2, \mathbf{s}'^{-jk}) \\ &\quad + (1 - p_m) E V_j(\omega_j', \xi_j', \mathbf{s}'^{-\mathbf{j}}) \Big] \Big\}, \end{aligned} \tag{10}$$

where

$$\zeta_j(\omega_j', \xi_j', \mathbf{s}'^{-j}, \omega_k', \xi_k', \mathbf{s}'^{-k}) = \frac{V_j(\omega_j', \xi_j', \mathbf{s}'^{-j})}{V_j(\omega_j', \xi_j', \mathbf{s}'^{-j}) + V_k(\omega_k', \xi_k', \mathbf{s}'^{-k})}. \tag{11}$$

As discussed earlier, we set a lower bound on the number of firms in the industry, denoted by $\underline{\mathrm{n}}$. When $n = \underline{\mathrm{n}}$, $p_m = 0$ and (10) reduces to (6).

In each period, sequence of events is the following.

1. Firms observe individual and industry states.
2. Pricing and investment decisions are made.
3. Profits and investment outcomes are realized.
4. Individual and industry states are updated before the mergers.
5. Mergers are drawn randomly (see Appendix for details).
6. The state variables of merged firms are adjusted and the industry states are updated accordingly.

2.3 Markov Perfect Equilibrium

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

1. $V(\omega, \xi, \mathbf{s})$ satisfies (10) and $x(\omega, \xi, \mathbf{s})$ is the optimal policy function;
2. $\pi(\omega, \xi, \mathbf{s})$ maximizes profits in the spot market;
3. $Q(\mathbf{s}', \mathbf{s})$ is the transition matrix that gives the probability of state $\mathbf{s}'$ given that the current state is $\mathbf{s}$.

Estimating the parameters in the static profit function—demand and supply parameters—is fairly standard. The one dynamic parameter in the model, 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), as defined above, from a starting value of c and use maximum likelihood estimation, as in Rust [1987] or Holmes and Schmitz [1995], 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 a similar approach in his study of the market for wide-bodied commercial aircraft. The main disadvantage of this approach is that the numerical solution for the MPE is computationally very intensive, despite some recent innovations by Pakes and McGuire [2001], Doraszelski and Judd [2004] and Weintraub et al. [2005].

The second is a two-step approach proposed by Bajari et al. [2006]. Their method allows for the estimation of the policy and value functions and for the recovery of structural parameters of the model without having to compute the MPE. They get around the problem of computing the equilibrium by assuming that the data we observe represent an MPE.[14] This assumption is not completely innocuous. For example, in the global auto industry firms often need to undergo structural changes with adjustments spread over many years. During this transition period, firms do not behave as they would in equilibrium. A prime example would be the three-year recovery plan Renault initiated at Nissan, when it took control of the troubled Japanese automaker. Similarly GM and Ford, having lost a big chunk of their market share in recent years, are undergoing massive restructuring. Their decisions during this restructuring phase are likely to be different from their decisions in a stable equilibrium. Nevertheless, the assumption that firms always play their equilibrium strategy sounds reasonable and, more importantly, does away with the need to compute the MPE.[15]

Given this assumption, the first step proposed by Bajari et al. [2006] is to estimate the state transition probabilities and the equilibrium policy functions directly from the observed information on investments and the evolution of the knowledge stock. Together with estimates of the static profit function, these can be used to obtain the value functions by forward-simulation. In the second step, the value function estimates from the first step are combined with equilibrium conditions of the model to estimate the structural parameter(s). In Section 4 we elaborate on these steps in some detail and explain how we apply them to our model. However, first we describe our data set in the next section.

[14] Alternative approaches that avoid computing the MPE at each iteration include Aguirregabiria and Mira [2007] and Pakes et al. [2006].

[15] Another implication of this assumption is that one has to rule out the possibility of multiple equilibria or simply assume that multiple equilibria may exist but the firms only play one and the same equilibrium in all periods. For a review of the problem of multiple equilibria in these models and its possible solutions see Section 6 in Doraszelski and Pakes [2006].

3 Data

We choose our sample period to be 1980–2005. This period covers most of the consolidation that has taken place in the industry over the last few decades. We limit the sample to the largest 13 firms (in terms of unit sales) that are active in 2005. 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 Renault never obtained majority control. The industry has seen significant consolidation over the last 25 years and working back to 1980, these 13 groups emerged from 22 initially independent firms. Table 2 lists different firms in the initial and the final year of the sample. In 2005, our sample accounted for slightly more than 85% of global automobile production. The remaining 15% was produced by a large number of small firms. Since patenting activity of these small firms is negligible and other information on them is spotty, we ignore these fringe firms.[16]

[Table 2 approximately here]

In order to estimate the parameters we need data on the following five variables: gross addition to a firm's knowledge (x); knowledge stock (ω); product features unobservable to the econometrician (ξ); market share; and prices. We describe each of these in some detail.

Our measure of gross addition to a firm's knowledge is the number of patents issued to a firm in a calendar year.[17]. Patent information is taken from the database of the National University of Singapore (NUS), which covers the period 1975–2004. Its main advantage over the more widely used NBER's patent database is that coverage is extended beyond 1999. Since different subsidiaries of the same firm might file for patents, we searched the database using several variations of the names of each firm and manually scrolled through the results to make sure that all appropriate patents were included.

Our measure of knowledge stock of a firm is its 'patent stock'. Using the number of patents issued to each firm for the period 1975–2004, we construct the patent stock using the perpetual inventory method: $\omega_{t+1} = (1 - \delta)\omega_t + \tilde{x}_t$.[18] The sum of patents awarded to each firm from 1975 to 1979 gives the initial patent stock at the beginning of the year 1980. We then depreciate this

[16] The only sizeable firm in 1980 not included in our sample is 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 shortly after 1979.

[17] Patents are widely used as a measure of innovation output. Despite some well known problems, they remain the most reliable 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, patents statistics remain a unique resource for the analysis of the process of technical change."[p. 1701]

[18] The difference between this equation and the law of motion for the knowledge stock, equation (7), is that the firm plans to obtain x_t patents, but the randomness in the innovation process yields the observed $\tilde{x}_t = x_t + \epsilon$ new patents, with $E(\tilde{x}_t) = x_t$.

patent stock at a constant rate δ and add the patents issued in 1980 to obtain the patent stock at the beginning of 1981. Continuing this forward, we can construct the patent stock for each firm up to the year 2004.

By definition, we do not observe ξ's. However, we can estimate them as residuals from the demand equation. We describe the demand equation and how to get ξ's from it in Section 4.1.1.

Our measure of innovation (i.e. the number of patents issued) is only available at the firm level and is assumed to be useful for the firm's worldwide activities. On the output side, we assumed that each firm produces a single model. Our empirical counterpart to output is the number of vehicles produced worldwide by each firm and its affiliates. This information is obtained using the Ward's Automotive Yearbooks (various years), supplemented by information from the online data center of Automotive News for the most recent years. Market share of a firm is computed as the ratio of the number of vehicles produced by a firm to the total number of vehicles produced in the global market.

To construct prices for the 'composite models', we estimate a hedonic price regression with the log prices of all available models in the market as dependent variable and a host of vehicle characteristics as explanatory variables – see Goldberg and Verboven [2001] for an example for the European car market.[19] We include a full set of firm-year interaction dummies; their coefficients capture the relative price for each firm over the sample period. We used GM as the base firm. This log relative price is exactly what we need to estimate the demand equation in (12) below, where we also have to express market share and knowledge of each firm relative to GM. For the current version of the paper we estimate the hedonic regression for the U.S. passenger vehicle market, updating the data set in Petrin [2002] to 2004. Figure 1 illustrates the evolution of these prices for a number of firms. We are in the process of compiling similar information for the European and Japanese markets to estimate prices representative of the firms' worldwide sales.[20]

[Figure 1 approximately here]

4 Estimation Methodology

In this section we describe our estimation methodology, which closely follows Bajari et al. [2006]. We add specific details where needed.

[19]Bajari and Benkard [2005] discuss the performance of hedonic pricing models when some product characteristics are unobservable.

[20]The relevant price will be the weighted average of the firm-year dummies from hedonic regressions in the three regions, using each firm's relative sales in the different regions as weight.

4.1 Step 1

In the first step, we estimate the demand and cost of production parameters, the transition probabilities and policy functions and use them to evaluate the value functions using forward simulations.

4.1.1 Estimation of Demand Parameters

The demand side in our model is static and we do not need the full model to estimate the demand parameters.[21] Following Berry [1994] we use (2) to write the log of the market share of firm j relative to a base firm 0 as

$$\log[\sigma_j(\cdot)/\sigma_0(\cdot)] = \theta_\omega \log[(\omega_j + 1)/(\omega_0 + 1)] + \theta_p \log[p_j/p_0] + [\xi_j - \xi_0]. \tag{12}$$

Using the observed market shares for σ_j and σ_0 and with data on ω's and prices, we can estimate the above equation by OLS to get the estimates for θ_ω and θ_p.[22] However, as producers use information about the unobserved vehicle quality (ξ) 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 upwardly biased. We use an IV estimator and follow the instrumenting strategy of Berry et al. [1995]: the sum of observable characteristics of rival products is an appropriate instrument for own price. In our case, this boils down to just the sum of ω's. The residuals from (12) are our empirical estimates of ξ in excess of the ξ for the GM.

4.1.2 Estimation of Production Cost Parameters

We do not observe marginal cost, but once we have the estimates for θ_ω, θ_p, and the vector ξ from the demand system, we use (4) to recover marginal costs.[23] Assuming firms are setting prices optimally, we solve for the marginal costs that rationalize the observed prices and market shares. We denote this new variable by mc. Assuming that mc is a linear function of ω, we can run a simple OLS regression of the form

$$mc_j = \gamma_0 + \gamma_1 \omega_j + \varepsilon_j. \tag{13}$$

We also use two alternative specifications: (i) we impose that $\gamma_1 = 0$, to study the case in which marginal cost is constant and does not vary with the technological knowledge of the firm; (ii) we use $\ln(\omega_j + 1)$ instead of ω_j, to impose diminishing returns to knowledge in terms of cost savings.

[21] We estimate the demand, cost, and policy functions pooling data across all years; time subscripts are omitted.

[22] Throughout this study we use GM as the base firm.

[23] While it is possible to impose optimal price setting in the previous step to increase precision, we chose not to as it would force the price elasticity to exceed unity.

4.1.3 Estimation of the Policy Function

The next task is to characterize the policy function from the 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 vector $\mathbf{s}$. Bajari and Hong [2005] discuss the resulting distribution of the estimator if this step is carried out nonparametrically. In our application, we observe at most 22 firms over 25 years, which forces us to estimate this relationship parametrically.

We postulate the following policy function:[24]

$$x_j = \alpha_0 + \alpha_1 \omega_j + \alpha_2 \omega_j^2 + \alpha_3 \sum_{k \neq j} \omega_k + \alpha_4 rank_j + e_j. \tag{14}$$

This equation postulates that the innovation a firm decides to carry out in any year is a concave function of the firm's own knowledge stock, the combined knowledge accumulated by its rivals in the industry, and the firm's rank in the industry in terms of knowledge stock. This allows for the firm's investments in additional knowledge to decline if the accumulated stock becomes large. It is even possible that investments are scaled back so much that its knowledge stock declines because of depreciation. A priori, we expect α_1 to be positive and α_2 to be negative. The sum of rivals' stocks and the firm's rank are included to allow each firm's investment to depend on the state of the industry and its own relative position in a parsimonious way.[25]

Since we have assumed that firms always play their equilibrium strategy, the OLS estimate of the above equation will in effect give us the equilibrium policy of the firm as a function of its own and the industry state. The error term in equation (14) captures the approximation error between the true policy function $x(\omega_j, \xi_j, \mathbf{s})$ and the one we estimate.

4.1.4 Estimation of the State Transition Function

The state transition function for the first state variable ω is given in (7). We will measure ω as the accumulated stock of past patents, which decreases exogenously in value because of economic obsolescence of knowledge and expiration of patent rights.[26] x will be measured as the number of

[24] We do not include ξ in this policy function. The reason is data constraint: we currently have around seven observations per firm for our demand estimation. These are not enough to estimate an AR(1) process for ξ for each firm separately. As we describe in Section 5.1.1 below, for our forward simulations we use average ξ's that are constant over time for each firm. We tried adding our empirical ξ's directly to the policy function but the coefficient was statistically insignificant and we had to lose some 70% of our data: we have 491 observations to estimate the policy function but just 155 observations on empirical ξ's. In view of these problems, we simply drop ξ from our policy function estimation. We hope to be able to overcome this problem when we have more data in future.

[25] Clearly, many alternative characterizations of the policy function are possible. We plan to investigate the robustness of our estimates to more flexible specifications.

[26] This idea of patent stock is similar to the one used by Cockburn and Griliches [1988].

new patents sought by the firm in the current year. The actual number of patents a firm will be awarded is $x+\epsilon$ and we assume that ϵ follows a normal distribution with zero mean and a standard deviation that is a positive function of x. The idea is to introduce a stochastic component into the next period's stock of knowledge.

The only parameters in the state transition function are the depreciation rate (δ) and the parameters of the distribution of ϵ. We do not estimate these parameters from the data, but assign them values that we believe are reasonable. We provide details in Section 5.1.4 below.

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

The law of motion for ξ that we specified in (8) poses further challenges. Ideally we would like to use the empirical ξ's as estimated in Section 4.1.1 above and estimate the law of motion as an AR(1) process for each firm separately. However, we do not have long enough firm-level time series data to estimate reliable AR(1) processes.[27] As an alternative, we compute the average of these empirical ξ's for each firm and use these averages as constant ξ's for each firm. Hence, in effect this second state variable is fixed over time. When two firms merge, we simply assign the average of individual firm ξ's to the merged firm.

4.1.5 Computation of the Value Function

We can put the previous four building blocks together to obtain a numerical estimate of the value function starting from any state $(\omega, \xi, \mathbf{s})$, given the structural parameter vector θ. For the simple model described above, 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 for this parameter (c_0). In step 2 below we use the equilibrium conditions of the model to derive a minimum distance estimator for this parameter.

To evaluate the value function we use forward simulation. We first explain our forward simulation procedure for the case when there are no mergers. Later we shall explain the additional steps to accommodate mergers. For an initial industry state $\mathbf{s}_0$, the estimated demand model and cost equation gives us the equilibrium period profit vector over all active firms in the industry according to (5).[28] The initial state directly determines optimal innovation, according to the estimated policy function, giving the net profit vector: $\pi_0(\mathbf{s}_0) - c_0 x_0(\mathbf{s}_0)$. Next, we use the state transition function to find the next period's state and denote it by $\mathbf{s}_1$. Following the same steps as above we can

[27]We have a total of 155 observations on prices. This gives slightly more than seven observations per firm.

[28]This subsumes the calculation of the equilibrium price vector by solving the system of first order conditions for all firms.

compute the expected net profit in period 1 as $\pi_1(\mathbf{s}_1) - c_0 x_1(\mathbf{s}_1)$. Since the value function is being evaluated at time 0, we need to discount the net profit in period 1 using an appropriate discount factor β. We continue this process for a sufficiently large number of periods T until the discount factor β^T gets arbitrarily close to zero. In other words, we evaluate the following equation:

$$V(\mathbf{s}_0|\theta) = \mathbb{E}\left[\sum_{t=0}^{T} \beta^t[(\pi_t(\mathbf{s}_t) - c_0 x_t(\mathbf{s}_t)]\right], \tag{15}$$

where the expectation is over future states. We run these forward simulations a large number of times, using different draws on ϵ_ω, the only source of uncertainty in the model, and take their average as a numerical estimate of $V(\mathbf{s}_0|\theta)$.[29]

We now explain how to forward simulate the value function in 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 (this value is calculated in the Appendix) they merge.[30] The ω of the merged firm is the sum of ω's of individual firms and ξ of the merged firm is the average of individual ξ's.[31] This changes the state of the industry in all calculations from that point onwards. In order to allocate the future profits of the merged firm to the value functions of each of the formerly independent firms, we need to calculate the share of each firm in their combined value should they have remained independent—according to equation (11). 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. As a result, the computational burden rises substantially. Starting from 22 firms, the first merger at time τ requires the calculation of additional value functions for 21 firms, although only $400 - \tau$ future periods will be considered, and so forth. We avoid this computational problem by first estimating the values without mergers. We then assign a share to each firm according to its relative value without mergers. For example, if firms A and B merge in period τ and their estimated values in no-merger case are $V_A(\omega_A, \mathbf{s})$ and $V_B(\omega_B, \mathbf{s})$ then, for each period after τ, we split the returns between them according to the following weights:

$$\zeta_A = \frac{V_A(\omega_A, \mathbf{s})}{V_A(\omega_A, \mathbf{s}) + V_B(\omega_B, \mathbf{s})}, \tag{16}$$

[29]The ξ_j/ξ_0 errors in equation (12) are fixed at the values that correspond to the year of $\mathbf{s_0}$. The ε and e errors in, respectively, equations (13) and (14), are fixed at zero throughout the simulations.

[30]If only one firm draws a merger, no merger takes place. If two firms draw a merger, they merge. If three firms draw a merger, we merge two of them randomly and the third remains unmerged. If four firms draw a merger, we merge two of them randomly and then merge the remaining two. And a similar procedure is used if more than four firms draw a merger. The probability to draw a merger (i.e. p) is different from the probability of actually experiencing a merger (i.e. p_m). We further clarify this distinction in the Appendix.

[31]The values of ξ (estimated as residuals from equation 12) for Chrysler and Daimler-Benz a year before their merger were -0.47 and 0.57. Two years after the merger, the ξ for the joint firm was 0.04. These estimates of ξ before and after merger support our assumption that the ξ of a merged firm is average of ξ's of the individual firms before the merger.

and

$$\zeta_B = \frac{V_B(\omega_B, \mathbf{s})}{V_A(\omega_A, \mathbf{s}) + V_B(\omega_B, \mathbf{s})}. \tag{17}$$

4.2 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(s) of the model contained in θ. The following steps assume that the model is identified and there is a unique true parameter vector θ_0. Bajari et al. [2006] 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 $\mathbf{x}'(\mathbf{s})$

$$V_i(\mathbf{s}, \mathbf{x}(\mathbf{s}), \theta) \geq V_i(\mathbf{s}, \mathbf{x}'(\mathbf{s}), \theta), \tag{18}$$

where $\mathbf{x}' \neq \mathbf{x}$ at the ith element, provided that we use the true value of the parameter vector $\theta_{\mathbf{0}}$.[32]

The minimum distance estimator for $\theta_{\mathbf{0}}$ is constructed as follows. For each firm i and state $\mathbf{s}$ we observe in the sample, we use the forward simulation method to calculate $V_i(\mathbf{s}, \mathbf{x}(\mathbf{s}), \theta)$. We do the same calculations using a number of alternative policy profiles $\mathbf{x}'(\mathbf{s})$ and compute the difference $V_i(\mathbf{s}, \mathbf{x}(\mathbf{s}), \theta) - V_i(\mathbf{s}, \mathbf{x}'(\mathbf{s}), \theta)$. We denote this difference by $d(i, \mathbf{s}, \mathbf{x}'|\theta)$. We then find d for all i, $\mathbf{s}$ and $\mathbf{x}'(\mathbf{s})$ for a given value of θ and compute the sum of the squared $\min\{d(i, \mathbf{s}, \mathbf{x}'|\theta), 0\}$ terms. This only penalizes the objective function if the alternative policy $\mathbf{x}'$ leads to a higher value function, which should not happen if $\mathbf{x}$ is the MPE profile. The θ with the smallest sum is our estimate of θ_0, i.e.

$$\hat{\theta}_{\mathbf{0}} = \arg\min \sum_{i,\mathbf{s},\mathbf{x}'} \Big[\min\{d(i, \mathbf{s}, \mathbf{x}'|\theta), 0\} \Big]^2. \tag{19}$$

5 Estimation Results

Here we present the estimation results; the sub-sections run parallel to Section 4 for easy reference.

5.1 Step 1

First, we present the estimates for the demand and (marginal) cost of production parameters, and the parameters in the policy function as observed in the data. We also indicate the parameters used to evaluate the state transitions and in the forward simulation of the value function.

[32] The Nash equilibrium assumes that all other firms play their equilibrium strategies, so we only change one element in the $\mathbf{x}$ vector, corresponding to firm i. Details on how we choose $\mathbf{x}'$ are in Section 5.2 below.

5.1.1 Estimates of Demand Parameters

The demand parameters using three different estimation methods—OLS, IV without time fixed-effects, and IV with time fixed-effects—are in Table 3. The impact of knowledge, θ_ω, is estimated similarly under all three specifications. The effect of knowledge on sales is positive, as expected, and estimated very precisely.

The price coefficient is estimated negatively, even with least squares (OLS) which ignores the correlation between price and unobserved product characteristics. Often, OLS estimates of discrete choice demand systems at the product level produce a price coefficient that is positive or very small in magnitude. A firm's patent stock seems to control for a lot of the usually unobserved quality variation that can lead to an upward bias in the price coefficient. Note that we do not observe any other characteristic; the price variable we use controls for any observable differences in the products offered by each firm.

Instrumenting for price with the knowledge stock of competitors leads to an estimated demand curve that is more elastic; consistent with previous results for the automotive industry (see for example Berry et al. [1995], Table III). Including time dummies increases the elasticity further. Using the estimates in the third column, the price elasticity varies between -7.2 and -1.96. Without enforcing the first order conditions for optimal price setting, all firms are estimated to price on the elastic portion of demand, consistent with the theory.[33] The price-marginal cost markups implied by these estimates are substantially lower than those obtained in other studies that estimate demand systems for car models. This is reasonable because we work at a much higher level of aggregation. At the firm level, a much larger fraction of costs will be variable than at the model level.

[Table 3 approximately here]

For the results in the third column, a 1% price increase has the same effect on demand as a 13% increase in the knowledge stock. GM is the base firm and all coefficients are identified off the variation relative to the GM values. A value of 0.20 for θ_ω means that a firm with a patent stock that is half of GM's patent stock, will on average have a market share that is 13% lower than GM's, holding price constant.

The average of ξ's for each firm are reported in Table 4. These ξ's are residuals from the demand equation, averaged for each firm. If a firm did not register most of its patents at the US Patent Office (USPO), then according to our calculations it is likely to have a low stock of knowledge, which may not explain a large part of the firm's market share. For such firms we would expect

[33] If we used the results in the first column, 10% of the observations in the sample would represent firms that price on the inelastic portion of demand, inconsistent with the theory. The elasticity of demand in the logit model can be calculated as $\theta_p p_j(1 - s_j)$, where s_j is the market share of firm j.

a large and positive ξ. The best examples are Fiat and Peugeot-Citreon. These big European firms have a sizable share in the global market but they have very few patents registered at USPO. Hence, as expected, we estimate positive and high ξ's for these firms. On the other extreme, there are firms that patent a lot but do not enjoy market shares commensurate with their patent stocks. For such firms we expect a large negative ξ. Honda and Chrysler are two such examples. Both these firms have been patenting a lot but their market shares relative to GM, Toyota and Ford are not as high as their patent stocks. Hence we expected large negative ξ's for these firms. The ξ's reported in Table 4 confirm this expectation.

5.1.2 Estimates of Cost of Production Parameters

We use the system of first order conditions, an equation like (4), for each active firm to recover its marginal cost. These are plotted against the patent stock in Figure 2. For low values of the patent stock there is no clear relationship between the two, but for a stock of a thousand or more patents, there is a clear negative relationship: a higher patent stock is associated with lower marginal cost. The stock of knowledge not only boosts demand, as we saw previously, it also reduces production costs. In this industry high patent stocks seem to be associated with product as well as process innovations.

[Figure 2 approximately here]

While the marginal costs we recover vary by firm and year, we need to be able to predict marginal costs at any possible state in order to estimate the model. Therefore, we run a simple OLS regression of the marginal cost on the patent stock, which gives the following results:

$$\begin{array}{cccccc} mc_{jt} & = & 0.9 & - & (7.36 \times 10^{-5}) \cdot \omega_{jt}, \\ & & (0.03) & & (2.05 \times 10^{-5}) \end{array}$$

Standard errors are in parentheses. When we assume that marginal costs are constant across firms and over time, the estimated marginal cost is 0.83. Finally, the estimated cost function when marginal cost is assumed to be a function of $\log(\omega)$ is:

$$\begin{array}{cccccc} mc_{jt} & = & 1.01 & - & 0.03 \log(\omega_{jt} + 1), \\ & & (0.08) & & (0.01) \end{array}$$

5.1.3 Estimates of the Policy Function

The policy function estimates are in Table 5. Results in the first column are OLS estimates for the entire sample. The coefficients on 'own stock' and 'own stock squared' have the expected positive and negative signs respectively. However, the coefficient on the squared variable is statistically insignificant. It turns out that this result is driven by a small number of outliers. To see this, we plot the number of new patents granted against the depreciated patent stock in Figure 3. When the patent stock is below 1200, there is a clear positive relationship between the flow of new innovations and the existing stock. When the patent stock exceeds 1200, there is no clear pattern, but the relationship is certainly not as steep as for the bottom-left observations. However, due to six outliers for which a firm is awarded more than 600 new patents in a single year, the relationship appears to be positive.

[Figure 3 approximately here]

Running the same OLS regression excluding these outliers, the estimates of both the linear and squared coefficients on a firm's own patent stock are higher (in absolute value), and estimated much more accurately. The R^2 also jumps from the first to the second column. The results now suggest a clear concave relationship between innovation and patent stock. When the patent stock of a firm is low, any increase in patent stock induces the firm to patent more. However, when the patent stock is sufficiently high, a further increase in the patent stock leads to less innovation.

[Table 5 approximately here]

The combined patent stock of a firm's rivals, on the other hand, is a negative predictor for innovation. The effect is estimated relatively precisely and large in absolute value. At the margin, a firm with a stock of one thousand patents, is estimated to increase its new patent target by 0.089 for each additional patent it holds. If all its 22 rivals held one additional patent, the firm is estimated to lower its target for new patents by 0.055. In fact, raising the patent stock for all firms in the industry by one is estimated to lower innovation for all firms that hold more than 1320 patents, which is not an uncommonly large stock as can be seen from Figure 3.

This result plays an important role when we simulate the value function forward. It restricts the number of patents a firm wants to accumulate over time, reinforcing the effect of the negative coefficient on the square of own patents. Without these effects the total patent stock in the industry would grow implausibly large, especially if we allow for process innovations that lower marginal costs.

Finally, the coefficient on own rank—a variable indicating how many other firms hold a larger patent stock—is not estimated significantly different from zero whether the regression is run on the full sample or omitting the outliers. Excluding the rank variable, results are in column 3 of Table 3, the estimated values for the other three coefficients hardly change, but their standard errors decline. The rank of a firm does not seem to play an important role in determining the number of patents it obtains; possibly because it does not contain much information once a firm's own patent stock and the combined stock of its competitors is controlled for.

5.1.4 Parameters of the State Transition Function

Two additional parameters affect the state transition function: the depreciation rate for the R&D stock (δ) and the standard deviation for the random shock to innovation (σ_ϵ). We do not estimate these parameters, but assign them reasonable values.

We assume a depreciation rate of 15% (i.e. $\delta = 0.15$).[34] This is the same depreciation rate that Cockburn and Griliches [1988] use to construct their R&D stock. This captures both patent expirations and the economic obsolescence of older knowledge. We set the standard deviation of ϵ_ω equal to 10% of x. As ϵ is normally distributed with mean zero, a firm that chooses to have 100 new patents in a particular year has a 67% probability of ending up with 90 to 110 patents.[35]

5.1.5 Computation of the Value Function

We assume that the number of consumers (m) in the industry follows a linear trend over time. We choose the slope of this trend line to match the data. Specifically we assume $m_{1980} = 37.4$ million and increase it by 1.1 million every year. The only other parameter we need is the discount rate β, which we set at 0.92. Given the (static) parameter estimates, an initial state of the industry, and a starting value for the dynamic parameter c_0, we simulate forward the evolution of the industry state and calculate profits for all firms as we go along. We simulate forward for 400 periods—$\beta^{400} = 3.27 \times 10^{-15}$—and construct the value functions as the present discounted value of the profit streams.

In the current setup, there is no permanent heterogeneity among the knowledge stocks of firms.[36] As a result, initial differences in knowledge stocks persist for some time but eventually vanish. The industry converges to a situation where some firms achieve a high and stable knowledge stock, only differing period by period because of random shocks to innovation. A high beginning patent stock gives the firms higher profits early on, which does lead to large differences in estimated value functions. Eventually, firms reach a patent stock at which marginal relative benefits on the demand side are outweighed by the cost of R&D, as captured by the concave shape of the value function and the negative effect of rivals' patent stock. Other firms have such a low initial patent stock that in the policy function the negative effect of rivals' knowledge dominates. These firms eventually cease innovating and their patent stocks depreciate away.

[34] Of course, we use the same value of δ to construct the patent stock from patent data.

[35] Although we choose the standard deviation of ϵ_ω arbitrarily, our estimate of the structural parameter c is not affected by this choice.

[36] ξ's are permanently different among firms but they do not appear in the estimated policy function.

5.2 Step 2

The only dynamic parameter in the model is the (average) R&D cost of a obtaining a new patent (c). Each period, a firm chooses the number of patents it would like to add to its knowledge stock, denoted by x. The R&D expenditure this requires upfront is $c \cdot x$ and the firm will obtain $x + \epsilon$ new patents by the next period. The minimum distance estimator as defined in (19) gives us an estimate of c.

To evaluate the objective function we have to specify alternative policies that differ from the equilibrium Nash policy. For each firm j and each industry state $\mathbf{s}$ we include 10 observations in the objective function by choosing $\mathbf{x}'(\mathbf{s}) = (\iota + ae_j)'\mathbf{x}(\mathbf{s})$, where ι is a vector of ones, e_j a vector of zeroes with a single one at position j (both vectors are of length n) and $a \in \{-0.10, -0.08, ..., -0.02, 0.02, ..., 0.08, 0.10\}$.

It turns out that the specification of the marginal cost function is crucial for the estimate of c. Assuming that the marginal cost is a linear function of the patent stock, as in equation (13), leads to an estimate of c that is 29.9 million dollars.[37] This specification makes patents very valuable and to fit the observed rate of patenting, the model estimates a very high R&D cost per patent.

In contrast, assuming a constant marginal cost the estimate falls to 11.6 million dollars. Including the logarithm of the patent stock instead of the linear term is likely to lead to an intermediate estimate. As a comparison, the median R&D-per-patent ratio we observe in the sample is 12.6 million dollars; very close to the c estimate if we assume a constant marginal cost.[38]

6 Findings

Thus far, we have estimated the structural parameters of the model which gave us some insights into the importance of innovation in the automotive industry. We found that innovation produces product and process innovation which affect both the demand and cost sides in plausible ways. Firms' optimal innovation policy depends on the state of the industry and the model produces a plausible estimate for the R&D cost of new patents. In this section, we use our model to study the interaction between innovation and market structure. The section is divided into three subsections. In the first subsection, we compare some predictions of the model with data. In the second subsection we study the key question: how changes in market structure interact with innovation. However, the primary goal of firms and policy makers is not innovation *per se*. Instead, the firms are more interested in the value that innovation creates and policy makers are more concerned about the consumer welfare. Hence, in the third subsection we study how mergers, the

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

[38] R&D expenditures per patent are very dispersed, even after dropping European producers that patent less often in the U.S. the standard deviation is \$12.5m, and the ratio has been increasing over time, totalling \$18.10m in 2004.

main source of changes in market structure in our model, affect firm value and consumer utility.

6.1 Model Predictions and Data

In this subsection we confront some predictions from the model to data. First, we compare the predictions from the model with the observed values for the top three firms in the sample: GM, Ford, and Toyota. These results are in Table 9. The model predicts lower market share for GM in both periods. The market shares for Ford and Toyota are predicted well for 1981 but are over predicted for 2004. Prices are over predicted for both Ford and Toyota. The reason is that prices in our model depend on utility, which in turn depends on knowledge stock and ξ. Since Ford and Toyota have knowledge stocks comparable to GM and ξ's that are higher than GM, the model predicts a higher price for their vehicles relative to GM's.

[Table 9 approximately here]

More importantly, the number of new patents the firms are predicted to apply for are estimated not too far off. The largest discrepancy is for Ford in 2004, but the actual number of new patents the firm filed that year (134) is an outlier itself (the average for Ford over the sample period is 260). The model also predicts the evolution of industry state reasonably well.[39] In Figure 4 we plot the mean and the standard deviation of ω by year as in the data and as predicted by the model. The predicted means and standard deviations are smoother than those from the data but still predict the data fairly well. While the estimated valuations for the three firms seem highly unrealistic, their evolution over time is more reasonable and incorporating fixed legacy costs for GM and Ford, as indicated in equation (3), would adjust the levels.

What the model does capture well is that firms with a patent stock above a threshold have an incentive to innovate a lot and converge to the leading firms. The simulations predict firms to fall into two groups; each of the three firms in Table 9 fall into the high patent stock–high price group. In the observed patent stock data, we do see the leading firms converging to similar levels, see Figure 5. The coefficient of variation across the five largest firms falls from 0.61 initially to 0.20 in 2004. To achieve this outcome, firms with a lower patent stock initially innovate more rapidly to catch up with the leaders.

[Figure 5 approximately here]

[39]By industry state we mean $\mathbf{s}_\omega$ because $\mathbf{s}_\xi$ does not evolve endogenously.

6.2 Market Structure and Innovation

The question that how market structure affects innovation has been extensively studied in the literature and the evidence is mixed. Examining evidence on the relationship between market concentration and innovation, Cohen and Levin [1989] (p.1075) cite some studies that found a negative relationship, some others that found a positive relationship and still others that found an inverted-U relationship. However, they conclude that "the empirical results concerning how firm size and market structure relate to innovation are perhaps most accurately decribed as fragile."(p.1078) More recently, Aghion et al. [2005] derive an inverted-U relationship between competition and innovation in a general equilibrium setting and find empirical support for it in the manufacturing industries in the UK. In their model, competition is completely exogenous. An increase in competition affects innovation positively if firms in the industry are technologically very close to one another and negatively if they are technologically far apart. The net effect depends on the steady state distribution of industries by technology gap. In their model, however, innovation does not affect the level of competition. Can our model offer some insights into this relationship and help resolve some of the issues? This was the main motivation of this study. Having estimated the parameters of the model, we are now in a position to examine that how market structure and innovation are related in our estimated model of the global automobile industry.

Technically, the market structure of the global auto industry can most appropriately be classified as an oligopoly and has been so since the inception of the industry. However, the extent of competition faced by firms in this industry has changed over time. We refer to these changes in competition as changes in market structure. We use two definitions of competition: the inverse of the degree of concentration in the industry; and the ratio of marginal cost to price. Our measure of innovation is simply the R&D intensity, which is defined as the ratio of R&D expenditures to sales. In our model, competition and innovation are determined simultaneously and evolve together. At the same time, random mergers reduce the extent of competition exogenously.

To study this relationship we begin with the actual state of the global auto industry in 1980. We set $\xi_j = 0$.[40] Given the state of the industry, firms make their pricing and R&D investment decisions. We record these decisions and compute the statistics of interest for individual firms as well as for the industry as a whole. We then draw mergers randomly and update industry state based on firms' R&D investments and the outcome of merger draws. We continue this process until the number of firms is down to five.[41] Figure 6 plots the results based on a random merger sequence.[42] This merger sequence is reported in Table 7.

[40]Since ξ does not enter the policy function, setting $\xi = 0$ does not affect the results of this subsection.

[41]We set the lower limit at five to allow for the industry to become highly concentrated.

[42]Although we report the results based on a single merger sequence here, the over all conclusions remain exactly the same regardless of the sequence of mergers.

In Figure 6(a), x-axis measures firm level competition as measured by the ratio of marginal cost to price.[43] On y-axis we plot innovation intensity, which is defined as the ratio of R&D expenditures to sales. The relationship between competition and innovation turns out to be inverted-U shaped: as the degree of competition increases, firms increase their innovative activity. However, when competition becomes too intense, the innovative activity declines. In this example, the innovative activity is at its maximum when mc/p ratio is around 0.59. The same holds in Figure 6(c) when competition is measured by one minus the market share of the firm. In this case, the innovative activity seems to peak when a firm controls about 10% of the market. In Figure 6(b) we average the mc/p ratio and innovation intensity for all the firms in each time period. The inverted-U pattern still holds, though the positively sloping portion of the inverted-U is much longer than the negatively sloping one. Figure 6(d) is similar to Figure 6(b) except that competition is measured by one minus the Herfindhal's index.[44] Although these results are robust to different random sequences of mergers, they are not robust to different initial states of the industry. Before examining the effects of a change in the initial industry state on this relationship, it is instructive to dig a little deeper and try to understand the reasons for this inverted-U relationship.

In Figure 7 (see thick lines) we decompose various components that generated Figure 6(b). Since in our model, mergers make the market more concentrated over time, we plot these variables against time. Figure 7 shows that over time: the industry becomes more concentrated (7(a)); sales continue to grow along a linear trend (7(b)); aggregate R&D in the industry first increases and then flattens out (7(c)); and R&D intensity in the industry first increases and then decreases (7(d)). Industry becomes concentrated because of mergers. Sales increase along a linear trend because we assume market size to increase linearly. The trend in aggregate R&D owes its shape to concave policy function and mergers. As the industry evolves, smaller firms tend to innovate a lot. They do so because, due to the steep value function, marginal addition to knowledge adds a lot to their value. However, in our estimated model there are decreasing returns to knowledge. As firms grow in knowledge, benefits from further innovation decline. Hence they innovate less. Mergers tend to expedite this process of decreasing returns by causing discrete increases in the knowledge stock. To see this, we add dotted lines to Figure 7 showing the evolution of the industry in absence of mergers. The dotted line in Figure 7(c) shows that industry R&D was bound to flatten over time even in the absence of mergers. However, mergers hasten this process by discretely increasing the knowledge of merging firms.

Now we go back to robustness of the inverted-U relationship that we saw in Figure 6. Our discussion in the preceding paragraph suggests that the inverted-U depends on the initial state of the industry. In the experiment that generated Figure 6 we set the initial state of the industry to

[43] This is just one minus the Lerner's Index. The Lerner's Index is defined as $\frac{p-mc}{p}$.

[44] Herfindhal's Index $= \sum_{j=1}^{n} s_j^2$. Where s_j is the market share of firm j.

be the same as the actual state of the auto industry in 1980. As we mentioned in Section 3, in 1980 there were 22 firms in our sample and almost half of them were really small in terms of their knowledge stock. It was this initial industry state that resulted in a faster growth in R&D than in sales and hence innovation intensity increased in the earlier periods. Once these firms grew bigger and merged, the decreasing returns caused R&D to flatten while the sales continued to grow along their linear trend and hence R&D intensity fell. The initial increase and later decrease in research intensity gave rise to the inverted-U relationship that we observed in Figure 6. We now repeat the same experiment but this time set the initial industry state to be the actual state of the auto industry in 2004. By 2004, the 22 firms in our sample in 1980 had undergone nine mergers and their number had fallen to just 13. This is already a concentrated industry state. When we start from this initial state we observe that the relationship between competition and innovation is generally positive, i.e. as industry becomes more concentrated, innovation falls.[45] We plot the results of this experiment in Figures 8 and 9, which are comparable with Figures 6 and **??**. Figure 8 shows that now the relationship between competition and innovation is generally positive.[46] The reason, as Figure 9 shows, is that now the industry is already concentrated to start with. This makes R&D grows slower than sales and hence R&D intensity starts to decline from the very beginning.

To sum up, forward simulations based on our estimated model of global auto industry show that if the industry is too fragmented (as the auto industry was in 1980) then some consolidation is likely to increase innovative activity. However, as industry becomes more concentrated (as it is now), further consolidation will negatively affect the intensity of innovation in the industry. These results depend on the decreasing returns to knowledge stock reflected in the concave policy function that we estimated.

The main message from these results is that, in general, competition is good for innovation. However, if an innovation intensive industry gets too fragmented, some reduction in competition can increase innovation. Our results show that the global auto industry has already consolidated enough and any further big mergers (like the one recently proposed between GM and Nissan and Renault) would hurt the innovative activity.

6.3 Market Structure, Value and Utility

Innovation is not an end in itself. From a producer's point of view, innovation is desirable because it creates value by improving quality and by lowering cost. From a consumer's point of view, innovation is desirable because it adds to utility by improving quality and by reducing price. Hence,

[45] The random merger sequence used in this experiment is reported in Table 8

[46] Though at the firm level there is still a small negative component that mainly arises due to the presence of some small firms in the sample in 2004. In the early stage of evolution of the industry, these firms individually innovate more although the overall industry innovation declines.

the ultimate question from welfare point of view is that how changes in market structure, and the consequent changes in the innovative activity, affect firm value and consumer utility.

In this subsection we look for an answer to this question with the help of three specific examples. First, we consider the merger between Daimler-Benz and Chrysler that took place in 1997. This is an example of two medium-sized firms merging. Second, we consider the merger between Hyundai and Kia in 1998. This is an example of two relatively small firms merging. Finally we consider a counterfactual merger between Renault-Nissan and Ford. In this counterfactual example, we let two already very big firms to merge. As we shall see below, the effects of a merger on value and utility depend, to some extent, on the size of the firms involved in the merger.

6.3.1 Mergers and Value

First, we consider the merger between Daimler-Benz and Chrysler in 1997.[47] This experiment begins with the actual state of the industry (**s**) in 1997. We consider two scenarios. Under the first scenario, no merger takes place and we let industry state evolve as predicted by the model for 20 years. We compute values (sum of the expected discounted profits) for all firms in each period. These values are reported as thin lines in Figure 10. Under the second scenario, we merge Daimler-Benz and Chrysler in the first period (i.e 1997) and repeat the same steps as we did under the first scenario. The values under this scenario are reported as thick lines in Figure 10.

In this model, mergers are almost always value destroyers for merging firms. The reason is that the value function is very steep at low levels of knowledge but very flat after the knowledge has reached a threshold. Hence, in most of the cases, when two firms merge, their combined knowledge stock is less valuable than the sum of their individual knowledge stocks. We see this in Figure 10 for the merger in question. It is, however, not clear that what happens to the value of other firms in the industry after the merger. In this example, the aggregate value of all the firms in the industry (including the merging firms) increases. The reason is that when medium-sized firms merge into one, competition in the industry is reduced and this helps rivals to gain a greater market share and enjoy a higher markup. In this example, the addition to the value of rivals more than offsets the loss in the value of merging firms and the net effect is an increase in the aggregate industry value.

In the second experiment we merge Hyundai and Kia in 1998. The results are reported in Figure 11. Like the previous experiment, the value of the merging firms decline after the merger. However, unlike the previous case, this time the merger hardly has any effect on the aggregate industry value. However, as time passes, the model predicts a slightly lower aggregate value with merger than without it (see Figure 11(b)).

Finally we consider a counterfactual merger between Renault-Nissan and Ford in 2004. This

[47] In this and next two experiments, no other mergers are allowed except the ones under review.

merger is different from the first two because the firms involved in it are individually very big. We report the results of this experiment in Figure 12. The effects on firm and industry value are similar to the first experiment above. At the firm level, the merging firm loses value. At the industry level, aggregate value increases. However, a notable difference is in the magnitude of the increase. A comparison between Figures 10(b) and 12(b) indicates that in the latter case the increase in the aggregate value is more than twice the increase in the former case. The reason is that when two big firms merge, the industry becomes a lot more concentrated than when two medium-sized firm merge. This increase in concentration helps rival firms by reducing the competition they face and by increasing their market share. The merging firms themselves lose because of decreasing returns.

To sum up, mergers result is lost value for merging firms regardless of the size of the merger. However, the rivals benefit if the merging firms are sizable. This benefit increases with the size of merging firms.

6.3.2 Mergers and Utility

What about the effects of these changes in market structure and innovative activity on consumers' utility? To answer this question, we repeat the same three experiments but this time record average consumer utility in each period. Results are in Figure 13. In Figure 13(a) the merger between Daimler-Benz and Chrysler does not have a noticeable effect on utility right after the merger. However, after 20 periods or so, the model predicts that average consumer utility will be lower with merger than without it. In Figure 13(b) the merger between Hyundai and Kia increases average utility and this increase persists over time. In Figure 13(c) the merger between Renault-Nissan and Ford results in immediate increase in utility. However, over time the utility drops sharply and eventually the utility with merger is much lower than the utility without it.

To sum up, the bigger the size of the firms involved in a merger, the greater is the negative effect on consumers' utility. On the other hand, if two small firms merge, average consumer utility may actually increase. These findings further support the findings in Section 6.2. There we found that letting firms merge in an already concentrated industry was harmful for innovation. Here we find that if we let bigger firms merge, it will have negative effects on consumers' average utility.

Before we close this section two comments are in line. First, ξ's of the merging firms, also play an important role in determining the effect of a merger on consumers' utility. Specifically, if two firms with positive ξ's merger, the negative effect on utility is greater than if two firms with negative ξ's merge.[48] The reason is that the firms with positive ξ's are already doing good on their own. By forcing them to merge we impose decreasing returns on well performing firms. One the other hand, when two firms with low ξ's merge, their individual values are so low that a merger actually

[48]Recall that these positive and negative ξ's are with reference to the ξ of General Motors.

has a healthy (or at least less harmful) effect on consumers' utility. Second, although we report three specific mergers, the general conclusions hold in many other experiments that we conducted but whose results we did not report.

7 Concluding Remarks

We construct a dynamic industry model of global automobile industry that allows for random mergers but no entry or exit. We estimate the structural parameters of the model using a new method proposed by Bajari and Benkard [2005] to estimate dynamic games. We show usefulness of their method in terms of saving in computation time and highlight some of the problems faced by us while employing their methodology.

We then use the model to study the interaction between market structure and innovation. Our principal finding is that, in the global auto industry, the effects of market structure on innovation depend on the initial state of the industry. If the industry is very fragmented, an increase in concentration (brought about by random mergers in our model) promotes innovative activity. However, if the industry is already concentrated, a further increase in concentration will discourage innovation. A general implication of this result is that depending on the initial conditions, an industry may exhibit positive, negative or inverted-U relationship between market structure and innovation.

The reason for interest in the question that how market structure affects innovation is not innovation per se. Our ultimate interest is how these changes in innovation, brought about by changing market structure, affect the net value of firms and utility of consumers. In our model, mergers are always value destroyers for merging firms. The reason is decreasing returns to knowledge: due to concave value function, when two firms merge, their knowledge stocks add up but their values do not. However, depending on the size of merging firms, mergers can have positive or negative effects on aggregate industry value. Specifically, the bigger the size of the merging firms the better it is for rival firms' values and hence for the aggregate industry value.

The effects on the average utility of consumers also depend on the size of merging firms. If two big firms merge, the average utility eventually declines. However, a merger between smaller firms may increase the average utility over time.

The policy message of this study is clear. The global automobile industry has already consolidated enough. Any further consolidation will be harmful for innovation. However, mergers between small firms (or between a small and a medium-sized firm) may still be beneficial not only for innovation but also for consumer welfare.

Like any research work, there are many dimensions along which the present study can be improved. First, data set needs to be extended. We would especially like to have data to run hedonic price regressions for European and Japanese firms separately to have better price estimates. We would also like to have more data on patents by European firms. The European patent office would be a natural place to look for such data. These data are not yet publicly available. Second, given the important role played by ξ in our model, the assumption of a constant ξ over time is not very satisfactory though, given the data constraints, it is a natural starting point. However, once we have the extended data set, the law of motion for ξ needs to be estimated more carefully. Third, with better data and better ξ's, the policy function needs to be improved by incorporating ξ's into it. Fourth, due to the above three limitations, the model predictions about firm values are way off mark. We hope that once these limitations have been overcome, the value estimates will be much better in line with the data. Another improvement to bring estimated values closer to the observed values could be to introduce fixed costs as the third state variable.

The above limitations notwithstanding, we firmly believe that dynamic models of competition provide the best modeling environment yet to study how market structure and innovative activity are related. In line with earlier surveys, we also believe that this relationship is industry specific and any generalization may not be very useful for policy purposes. For example, if our findings are any guide, the anti-trust authorities in US, Europe and Japan should discourage any further mergers between big auto firms.

We also believe that arguments like whether perfect competition is better for innovation or monopoly may be of some theoretical interest but are not useful for any practical purposes. Take the example of the auto industry. It is impossible to imagine that in near future this industry will get any where closer to either perfect competition or monopoly. In fact, Klepper [2002b] and Klepper [2002a] make a convincing case for the US automobile industry that it has settled into an oligopoly. We expect a similar evolution for the global auto industry. Hence the arguments about whether perfect competition or monopoly is better for innovation are fruitless, at least as far as the auto industry is concerned. A much more relevant question is that given the current state of the industry whether further consolidation will help or discourage innovation. Such questions can best be analyzed within the framework used in this paper.

References

P. Aghion and R. Griffith. *Competition and Growth: Reconciling Theory and Evidence*. The MIT Press, Cambridge, MA, 2005.

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 C.-Y. Ho. A Dynamic Oligopoly Game of the US Airline Industry: Estimation and Policy Experiments. Working paper, University of Toronto, 2006.

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

S. Ahn. Competition, Innovation and Productivity Growth: A Review of Theory and Evidence. OECD, Economics Department Working Papers No.317, 2002.

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

P. Bajari and C. L. Benkard. Hedonic Price Indexes with Unobserved Product Characteristics, and Application to Personal Computers. *Journal of Business and Economic Statistics*, 23(1): 307–347, January 2005.

P. Bajari and H. Hong. Semiparametric Estimation of a Dynamic Game Model of Incomplete Information. Working Paper, 2005.

P. Bajari, C. L. Benkard, and J. Levin. Estimating Dynamic Models of Imperfect Competition. *Econometrica*, forthcoming, 2006.

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.

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

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, Amesterdam: North-Holland, 1989.

A. Collard-Wexler. Plant Turnover and Demand Fluctuations in the Ready-Mix Concrete Industry. Working paper, NYU Stern, 2005.

U. Doraszelski and K. L. Judd. Avoiding the Curse of Dimensionality in Dynamic Stochastic Games. Working paper, Harvard University, 2004.

U. Doraszelski and A. Pakes. A Framework for Applied Dynamic Analysis in I.O. Working paper, Harvard University, 2006.

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.

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.

T. J. Holmes and J. A. J. Schmitz. On the Turnover of Business Firms and Business Managers. *Journal of Political Economy*, 103(5):1005–1038, 1995.

M. I. Kamien and N. L. Schwartz. *Market Structure and Innovation*. Cambridge University Press, Cambridge, UK, 1982.

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

S. Klepper. Firm Survival and the Evolution of Oligopoly. *RAND Journal of Economics*, 33(1): 37–61, 2002b.

A. Pakes and P. McGuire. Stochastic Algorithms, Symmetric Markov-perfect Equilibrium and the Curse of Dimensionality. *Econometrica*, 69(5):1261–1281, 2001.

A. Pakes, M. Ostrovsky, and S. Berry. Simple Estimators for the Parameters of Discrete Dynamic Games (with entry/exit examples). *Review of Economic Studies*, forthcoming, 2006.

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

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

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

J. A. Schumpeter. *Capitalism, Socialism and Democracy*. Harper, New York, 3 edition, 1950.

A. Sweeting. Estimating the Cost of Format Switching. Working paper, Northwestern University, 2006.

G. Y. Weintraub, C. L. Benkard, and B. V. Roy. Markov Perfect Industry Dynamics with Many Firms. Working paper, Stanford University, 2005.

A Appendix: The Merger Technology

The purpose of this Appendix is to show how we can impute the probability that a firm will be involved in a merger (p_m), which is used in equation (10), from the observed number of mergers in the data.

Let there be $n > 1$ firms in the industry and let p be the probability (same for all firms) that a firm will be up for merger this period. This probability p differs from the merger probability p_m, because mergers only take place if at least two firms are up for merger.

With only two active firm, $p_m = p^2$. With more active firms, the probability that two firms merge at the end of the period 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 $\binom{n}{k}p^n(1-p)^{n-k}$. The last term in the above series of probabilities depends on whether n is an odd or an even number. If n is odd, the last term would be $\frac{n-1}{n}P_{n-1}^{n-1}$ and if n is even, the last term would simply be P_{n-1}^{n-1}. Adding these terms together, we can write the sum as

$$p_m(n,p) = \Big[\sum_{i\in O, i<n} P_i^{n-1} + \sum_{i\in E, i<n} \frac{i}{i+1} P_i^{n-1}\Big] \cdot p, \tag{20}$$

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. Our next task is to impute p from the data. In our sample we observe 9 mergers in 26 years. We shall assume for the sake of simplicity that these mergers are evenly spread over the entire sample period. Then the expected number of mergers in any period is $\frac{9}{26}$. We make this expected number of mergers depend on n in a simple way. Let $\overline{n}$ be the average number of firms in the industry per period. Then we may write the expected number of mergers as

$$E(M) = \frac{9}{26} \cdot \frac{n}{\overline{n}}. \tag{21}$$

If the actual number of firms in the industry is equal to the average over the sample period, we expect $\frac{9}{26}$ mergers to take place in that period (or we expect 1 merger in every $\frac{26}{9}$ such periods). If the actual number of firms is above the average we expect more mergers and vice versa. The simple idea behind (21) is that as n declines, we reduce the expected number of mergers that we want to match by choosing an appropriate p. In other words, we want to make p small as n declines.

So far we have been trying to get a reasonable number of expected mergers from the data. But if we know p and n, we can easily derive the expected number of mergers by using the following equation

$$E(M) = \sum_{M=1}^{\lfloor \frac{n}{2} \rfloor} M \cdot P(M), \tag{22}$$

where $P(M)$ is the probability that M mergers take place. If $M = 1$ then $P(M) = (P_2^n + P_3^n)$. In words, 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. If 3 firms are up for merger, we shall pick 2 at random and the third will not merge and remain independent. Equation (22) can explicitly be written as

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

The last equation implicitly defines p in terms of n and $E(M)$. But in (21) we have already defined $E(M)$ as a function of n. Hence, given n we can use (21) to solve for $E(M)$ and then use the last equation above to solve for p. Once we have p, we can use (20) to get p_m.

Table 1: R&D expenditures by industry in selected countries for 2003 (in PPP $m)

Industry	ISIC Rev. 3	OECD	USA	EU	Japan
Total business sector	15-99	577,316	204,004	125,591	84,676
Pharmaceuticals	24.2	57,541	15,962	16,850	6,363
Office, accounting, and computing mach.	30	24,136	7,664	2,568	10,764
Radio, television, telecommunications eq.	32	71,623	22,399	13,812	11,081
Medical, precision, optical instruments	33	38,440	20,400	6,782	3,619
Motor vehicles	34	76,199	17,034	21,258	12,765
		(13.2%)	(8.3%)	(16.9%)	(15.1%)
Aircraft and spacecraft	35.1	34,065	15,731	8,518	385
Wholesale & retail trade; repairs	50-52	30,066	26,580	1,039	578
Computer and related activities	72	38,189	19,854	6,990	1,814
R&D	73	22,417	12,460		4,951

Note: OECD is USA, EU, Japan, Canada, Korea, Norway, Poland, Czech Republic.

Source: OECD ANBERD database, Version 2, 2006.

Table 2: Firms in the sample with market share (worldwide unit sales)

1980		2005	
($n = 22$)	market share*	($n = 13$)	market share
BMW	0.010	BMW	0.020
Chrysler	0.036	Daimler-Chrysler	0.071
Daimler-Benz	0.020		
Fiat	0.047	Fiat	0.030
Ford	0.148		
Jaguar	0.001	Ford	0.115
Volvo	0.010		
Rover Group	0.014		
Daewoo	0.0004		
Saab	0.003	General Motors	0.123
General Motors	0.174		
Honda	0.028	Honda	0.050
Hyundai	0.002	Hyundai-Kia	0.055
Kia	0.001		
Mitsubishi	0.028	Mitsubishi	0.020
Peugeot (PSA)	0.044	Peugeot (PSA)	0.050
Nissan	0.072	Renault-Nissan	0.089
Renault	0.056		
Suzuki	0.017	Suzuki	0.031
Daihatsu	0.013	Toyota	0.122
Toyota	0.090		
Volkswagen	0.068	Volkswagen	0.077
Total	0.882	Total	0.853

Notes: *1982;

Source: Mostly Ward's Automotive Yearbook (various years)

Table 3: Demand Parameter Estimates

	OLS	IV	IV
θ_ω	0.22^{***}	0.20^{***}	0.20^{***}
	(0.03)	(0.04)	(0.04)
θ_p	-1.11^{***}	-1.82^{*}	-2.62^{**}
	(0.19)	(0.96)	(1.32)
Time Fixed-Effects	No	No	Yes
R^2	0.38	0.32	-
No. of observations	155	155	155

Table 4: Average ξ's

Firm Name	Average ξ	Firm Name	Average ξ	Firm Name	Average ξ
BMW	-0.06	Honda	-0.47	Renault	na
Chrysler	-0.50	Hyundai	-1.11	Renault-Nissan	na
Daewoo	na	Hyundai-Kia	na	Rover Group	na
Daihatsu	na	Jaguar	na	Saab	na
Daimler-Benz	0.34	Kia	na	Suzuki	-0.68
Daimler-Chrysler	0.16	Mitsubishi	-0.72	Toyota	0.43
Fiat	0.94	Nissan	-0.15	Volkswagon	0.73
Ford	0.55	Peugeot-Citroen	1.03	Volvo	na

Table 5: Policy Function Estimates

Independent Variable	OLS	OLS (without outliers)	OLS (without outliers)
Own Patent Stock	0.19^{***}	0.28^{***}	0.28^{***}
	(0.02)	(0.02)	(0.01)
Own Patent Stock Squared	-9.04×10^{-6}	$-5.55 \times 10^{-5***}$	$-5.51 \times 10^{-5***}$
	(7.38×10^{-6})	(7.41×10^{-6})	(6.21×10^{-6})
Sum of Rivals' Patent Stocks	$-1.64 \times 10^{-3*}$	$-2.48 \times 10^{-3***}$	$-2.46 \times 10^{-3***}$
	(9.39×10^{-4})	(7.54×10^{-4})	(7.33×10^{-4})
Own Rank	-1.60^{*}	0.07	–
	(0.97)	(0.80)	
R^2	0.80	0.84	0.84
No. of observations	491	485	485

Table 6: List of Abbreviations

Abrvn	Firm Name	Abrvn	Firm Name	Abrvn	Firm Name
BMW	BMW	GMS	General Motors	PSA	Peugeot-Citrion
CHR	Chrysler	HND	Honda	REN	Renault
DBZ	Daimler-Benz	HYK	Hyundai-Kia	ROV	Rover Group
DCH	Daimler-Chrysler	HYU	Hyundai	SAB	Saab
DSU	Daihatsu	JAG	Jaguar	SZK	Suzuki
DWO	Daewoo	KIA	Kia	TYO	Toyota
FIA	Fiat	MTB	Mitsubishi	VOL	Volvo
FRD	Ford	NSN	Nissan	VOW	Volkswagon

Table 7: Merger Sequence 1

Period	Merging Firms
2	CHR(139) NSN(674)
3	KIA(0) ROV(0)
4	DWO(0) FRD(443)
5	REN(12) $\boxed{\text{DWO FRD}}$
6	JAG(0) TYO(781)
7	VOL(3) $\boxed{\text{KIA ROV}}$
9	SAB(43) $\boxed{\text{JAG TYO}}$
11	SZK(6) $\boxed{\text{VOL KIA ROV}}$
13	FIA(0) $\boxed{\text{SZK VOL KIA ROV}}$
14	$\boxed{\text{REN DWO FRD}}$ $\boxed{\text{FIA SZK VOL KIA ROV}}$
18	HYU(0) $\boxed{\text{REN DWO FRD FIA SZK VOL KIA ROV}}$
20	VOW(179) $\boxed{\text{SAB JAG TYO}}$
24	DSU(7) GMS(1189)
27	$\boxed{\text{CHR NSN}}$ $\boxed{\text{VOW SAB JAG TYO}}$
50	PSA(0) MTB(199)
53	$\boxed{\text{HYU REN DWO FRD FIA SZK VOL KIA ROV}}$ $\boxed{\text{DSU GMS}}$
56	BMW(28) $\boxed{\text{HYU REN DWO FRD FIA SZK VOL KIA ROV DSU GMS}}$

Notes: (1) Numbers in parenthesis are the first period (i.e. 1980) knowledge stocks.

(2) Names in boxes are the firms that had merged before the current period .

(3) See Table 6 for the list of abbreviations.

Table 8: Merger Sequence 2

Period	Merging Firms
3	PSA(9) MTB(516)
7	BMW(396) $\boxed{\text{PSA MTB}}$
10	SZK(253) HND(3374)
16	FRD(1654) $\boxed{\text{SZK HND}}$
19	VOW(314) $\boxed{\text{FRD SZK HND}}$
34	FIA(47) $\boxed{\text{VOW FRD SZK HND}}$
35	TYO(2249) DCH(2619)
43	HYK(560) $\boxed{\text{BMW PSA MTB}}$

Notes: (1) Numbers in parenthesis are the first period (i.e. 2004) knowledge stocks.

(2) Names in boxes are the firms that had merged before the current period.

Table 9: Model predictions for top firms

	GM		Ford		Toyota	
	1981	2004	1981	2004	1981	2004
Market share (actual)*	0.174	0.121	0.115	0.098	0.090	0.115
Market share (predicted)	0.090	0.083	0.104	0.140	0.111	0.150
Price (actual)	1.000	1.000	0.858	0.970	0.901	0.961
Price (predicted)	1.000	1.000	1.080	1.074	1.055	1.024
New patents (actual)	328	381	149	134	241	405
New patents (predicted)	267	319	123	298	196	338
Patent stock (actual)	1189	1896	443	1654	781	2248
Value (predicted)	73	107	184	336	163	304

Notes: * 1982 instead of 1981;

Prices are relative to GM; valuations are in billion of dollars, assuming $10,000 per GM vehicle.

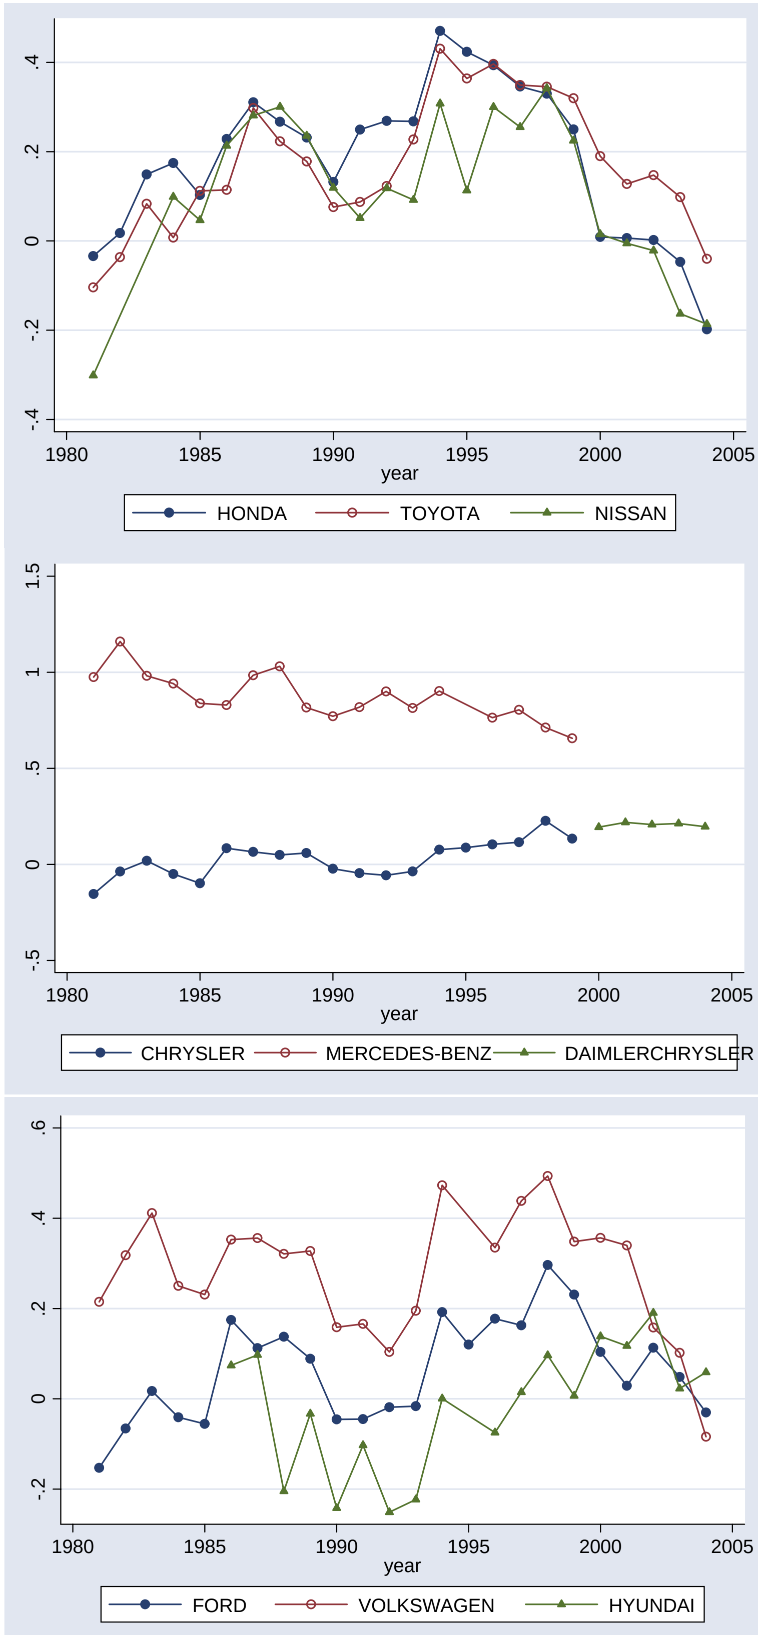

Figure 1: Hedonic prices for a number of firms (relative to GM)

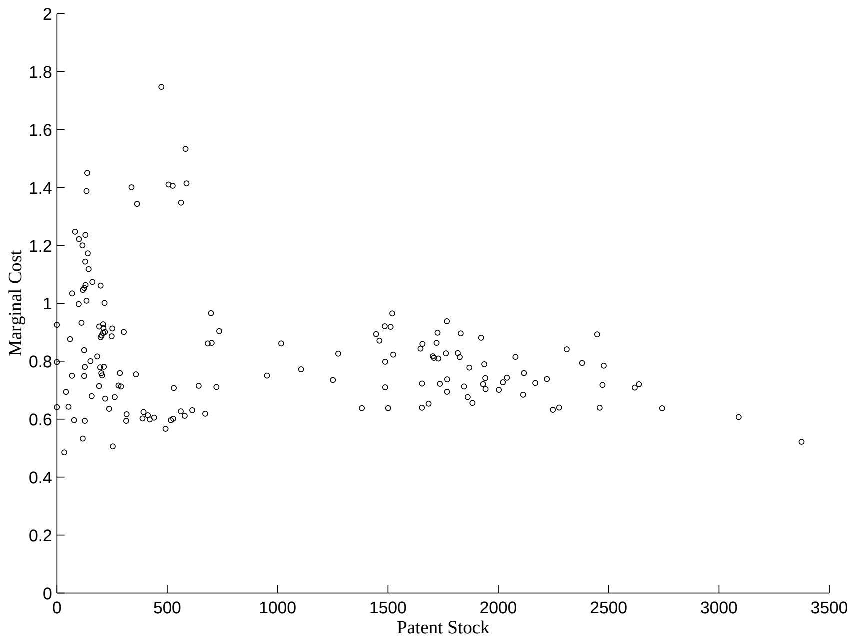

Figure 2: Marginal Cost and Patent Stock

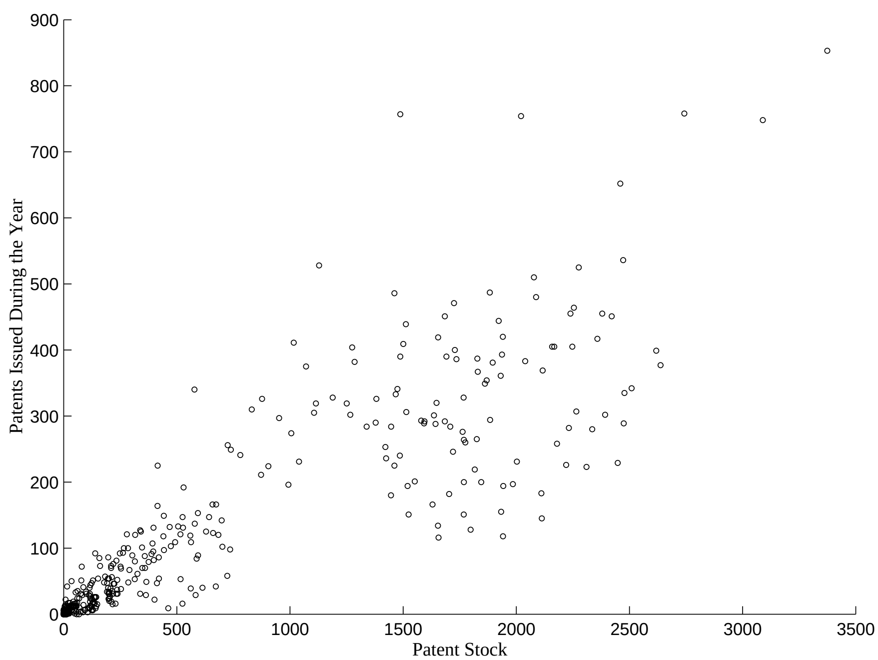

Figure 3: Patents and Patent Stock

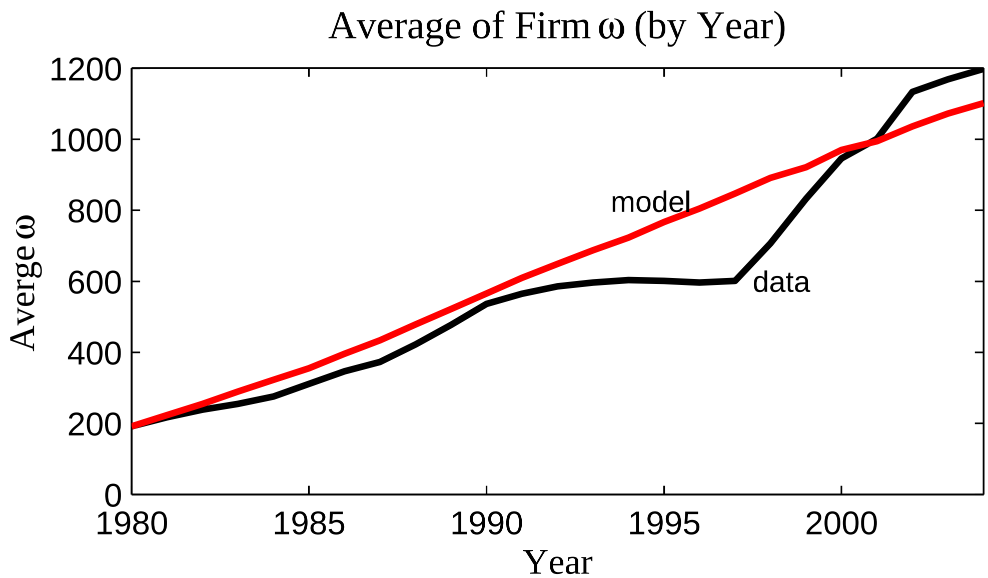

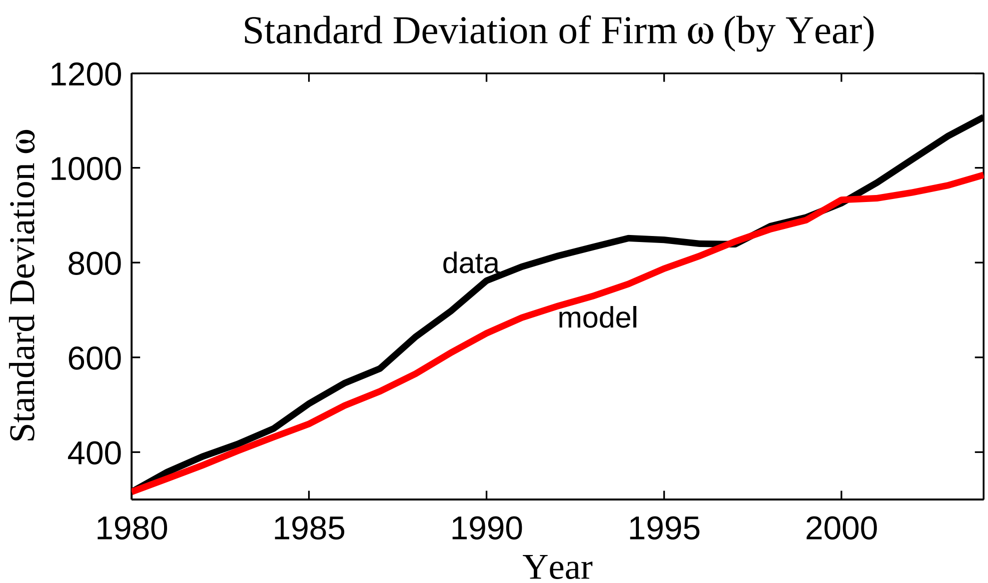

Figure 4: Evolution of ω's: Model (average of 30 random sequences of mergers) and Data

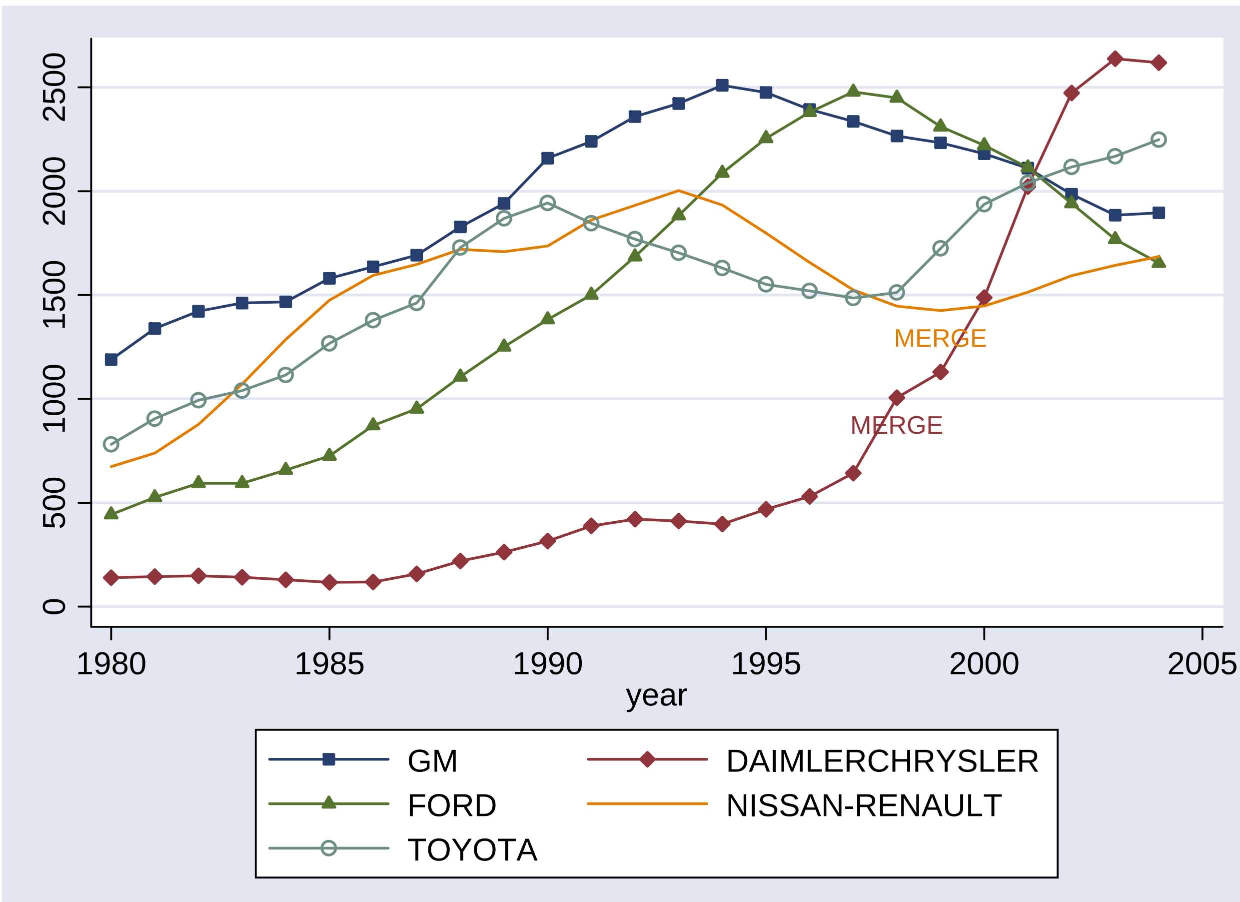

Figure 5: Patent stocks for the largest five firms converge

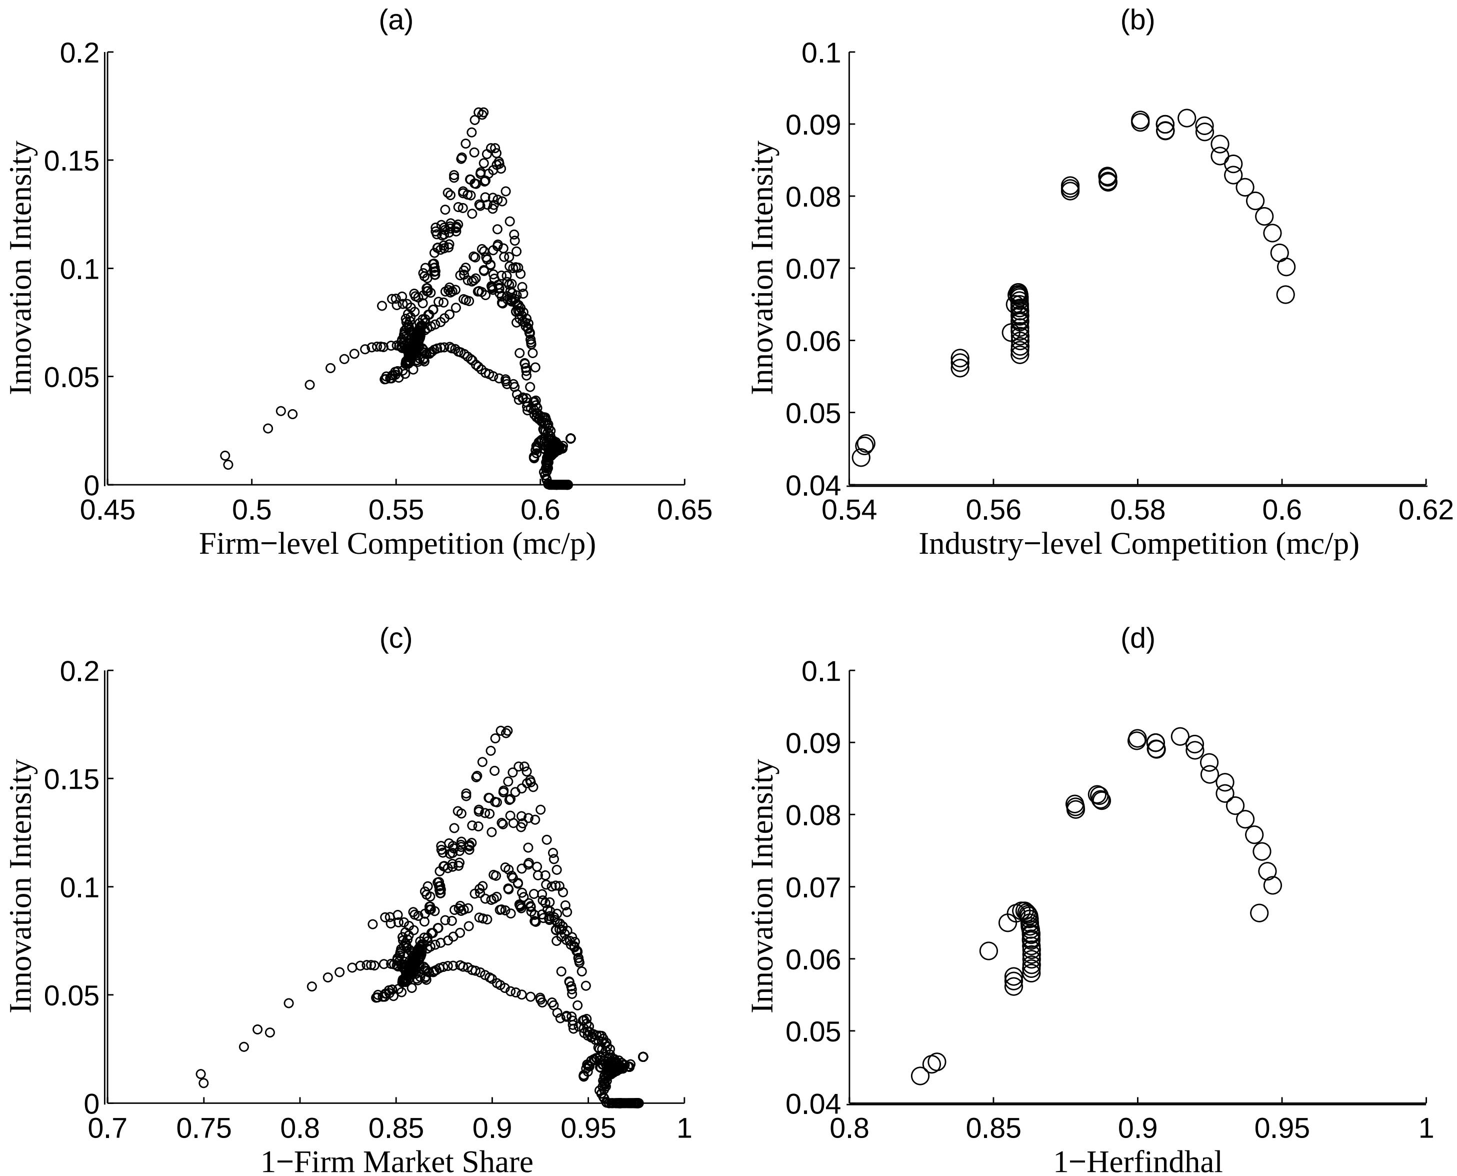

Figure 6: Competition and Innovation at the Firm and Industry Level

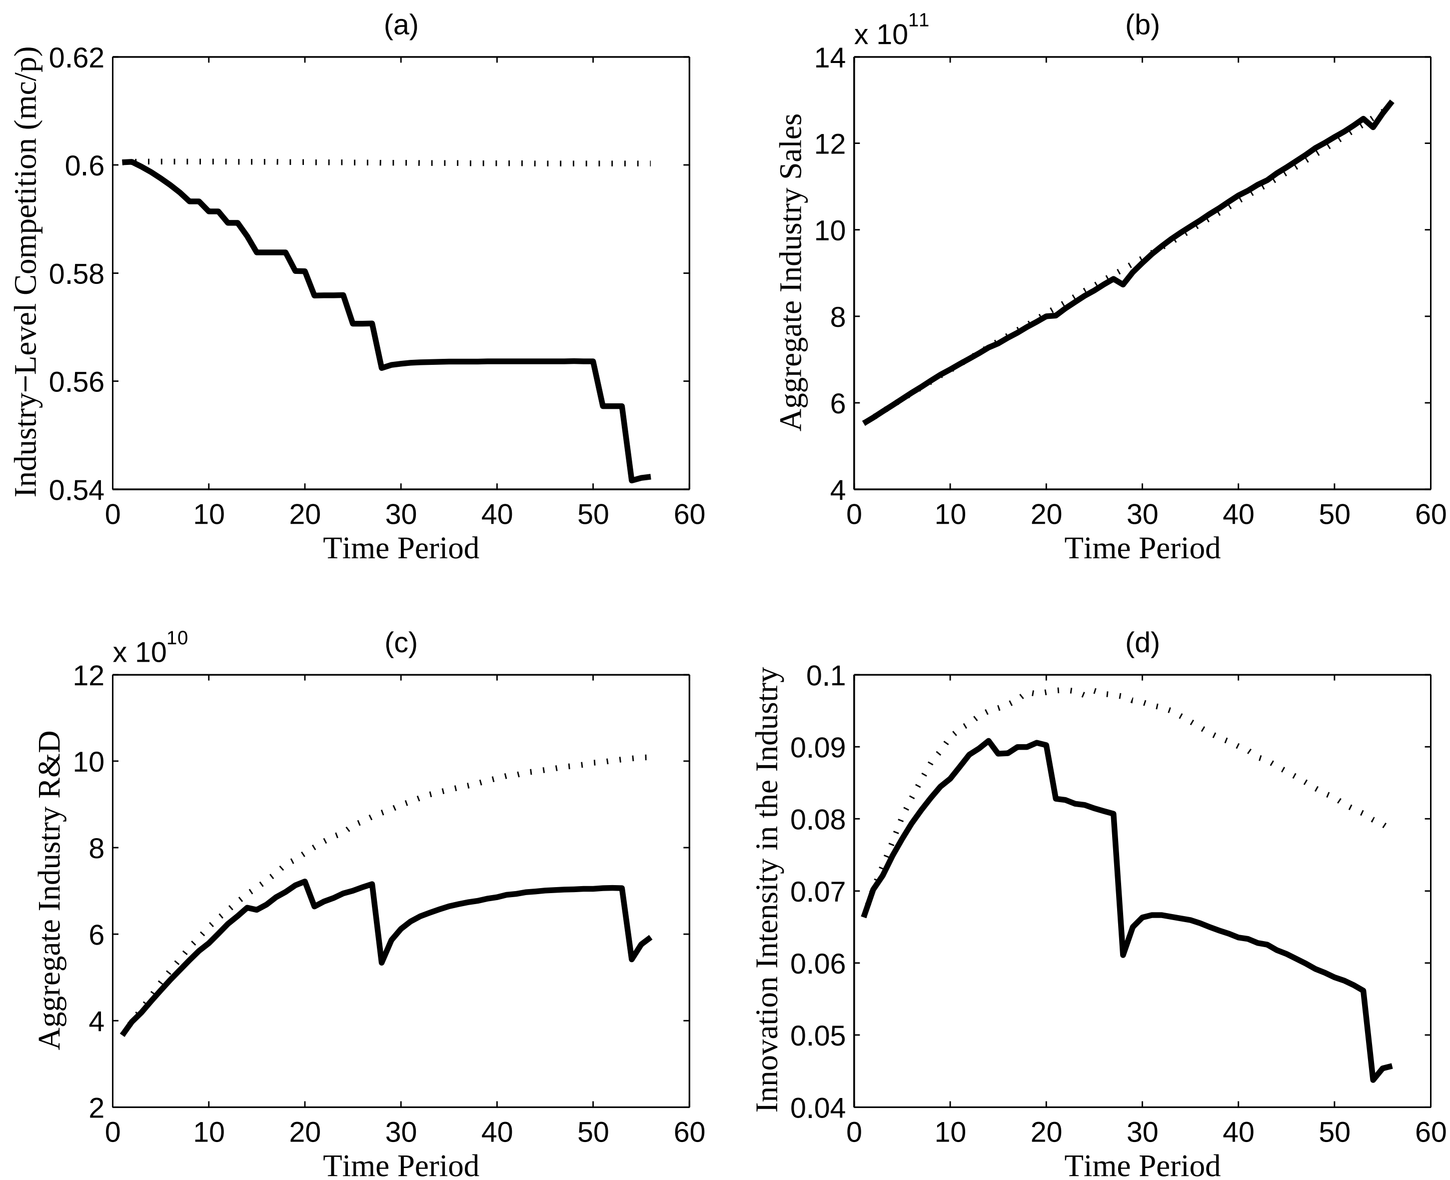

Figure 7: Competition and Innovation over Time, With and Without Mergers (Industry Level)

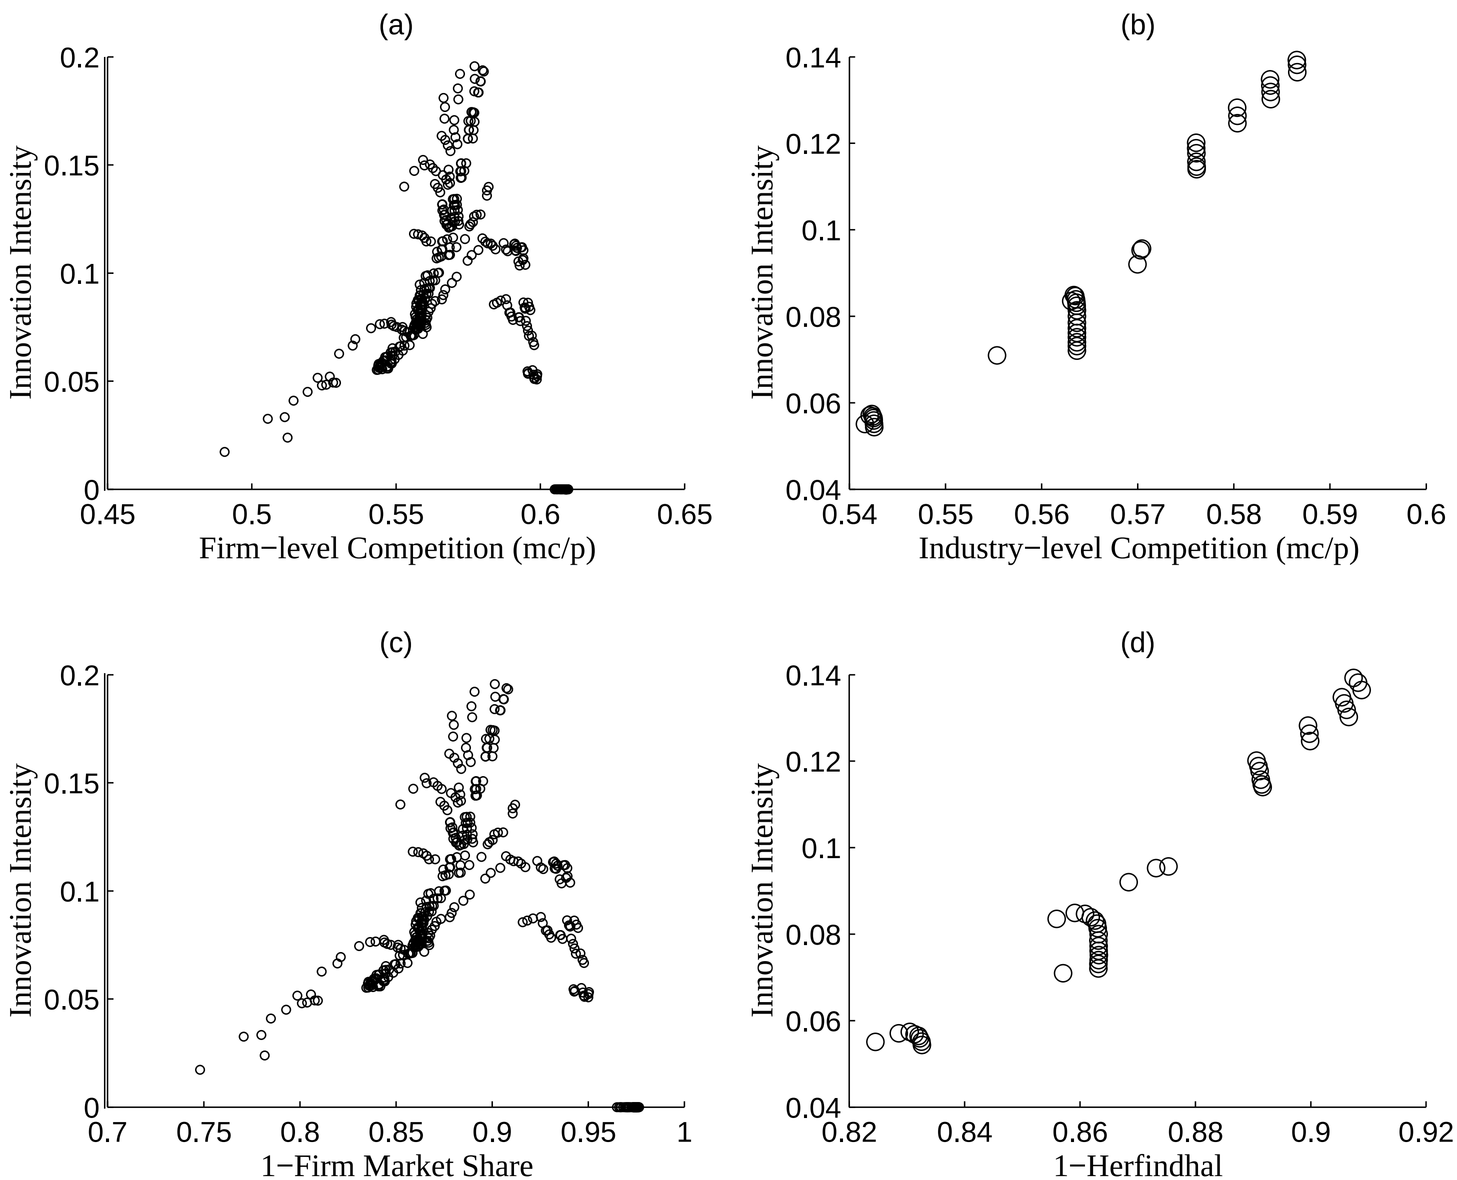

Figure 8: Competition and Innovation at the Firm and Industry Level when the Industry is Concentrated in the Initial State

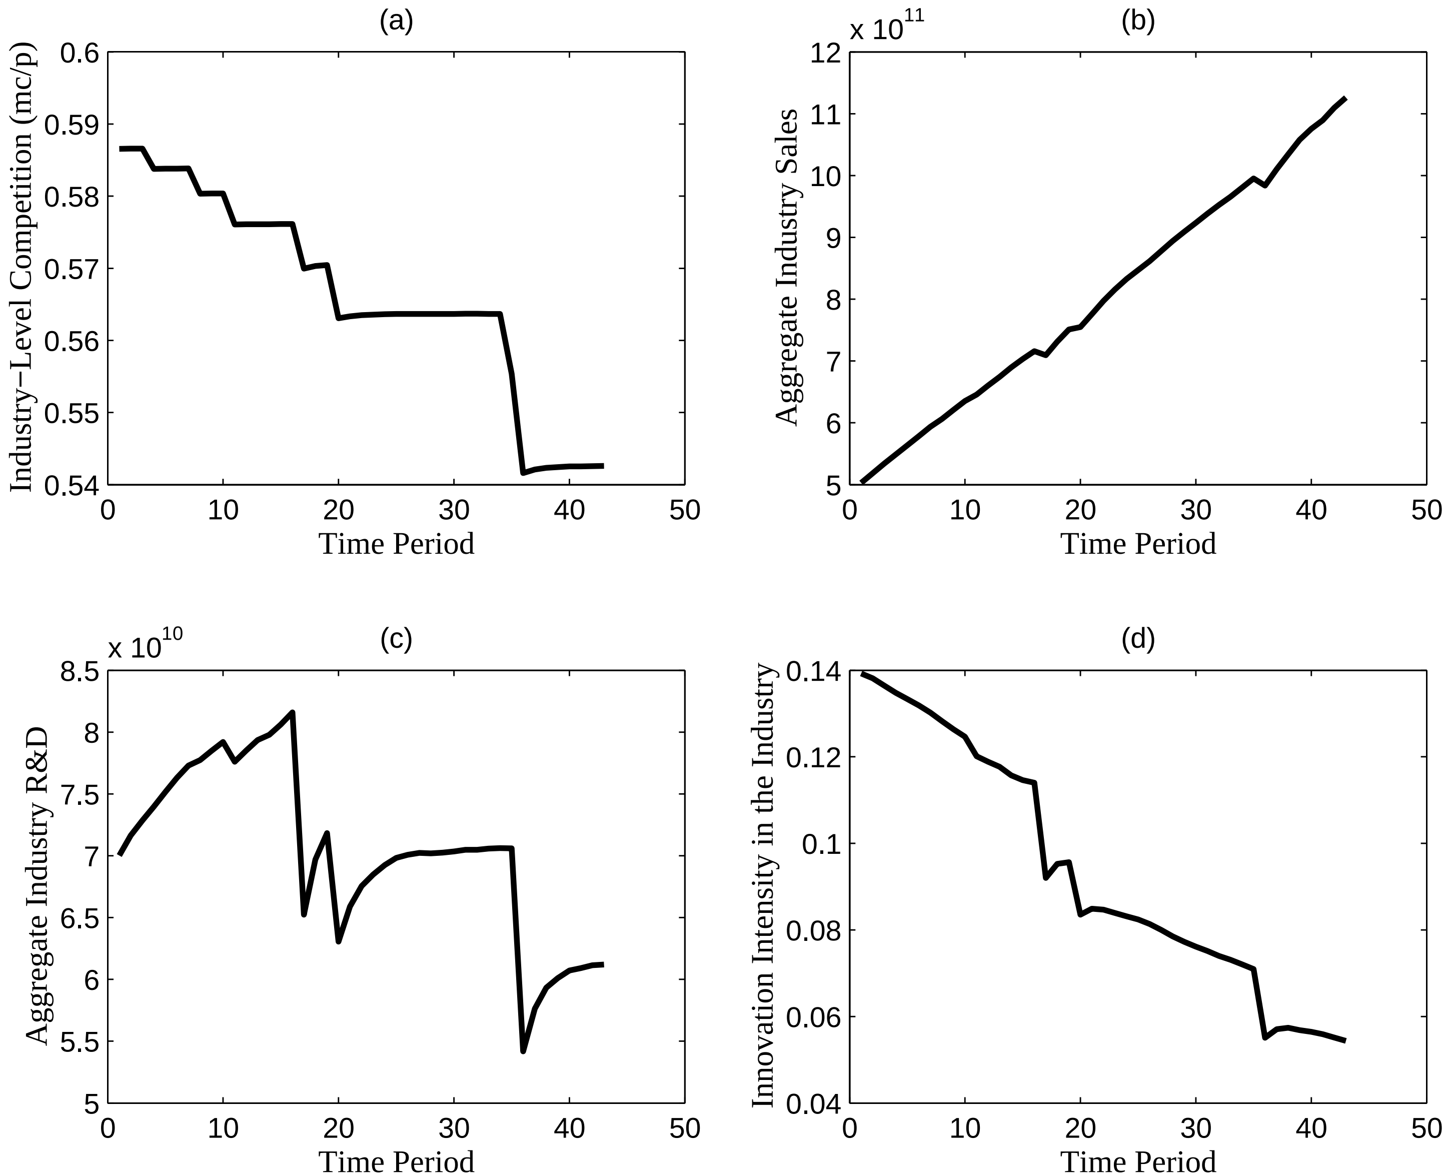

Figure 9: Competition and Innovation over Time (Industry Level)

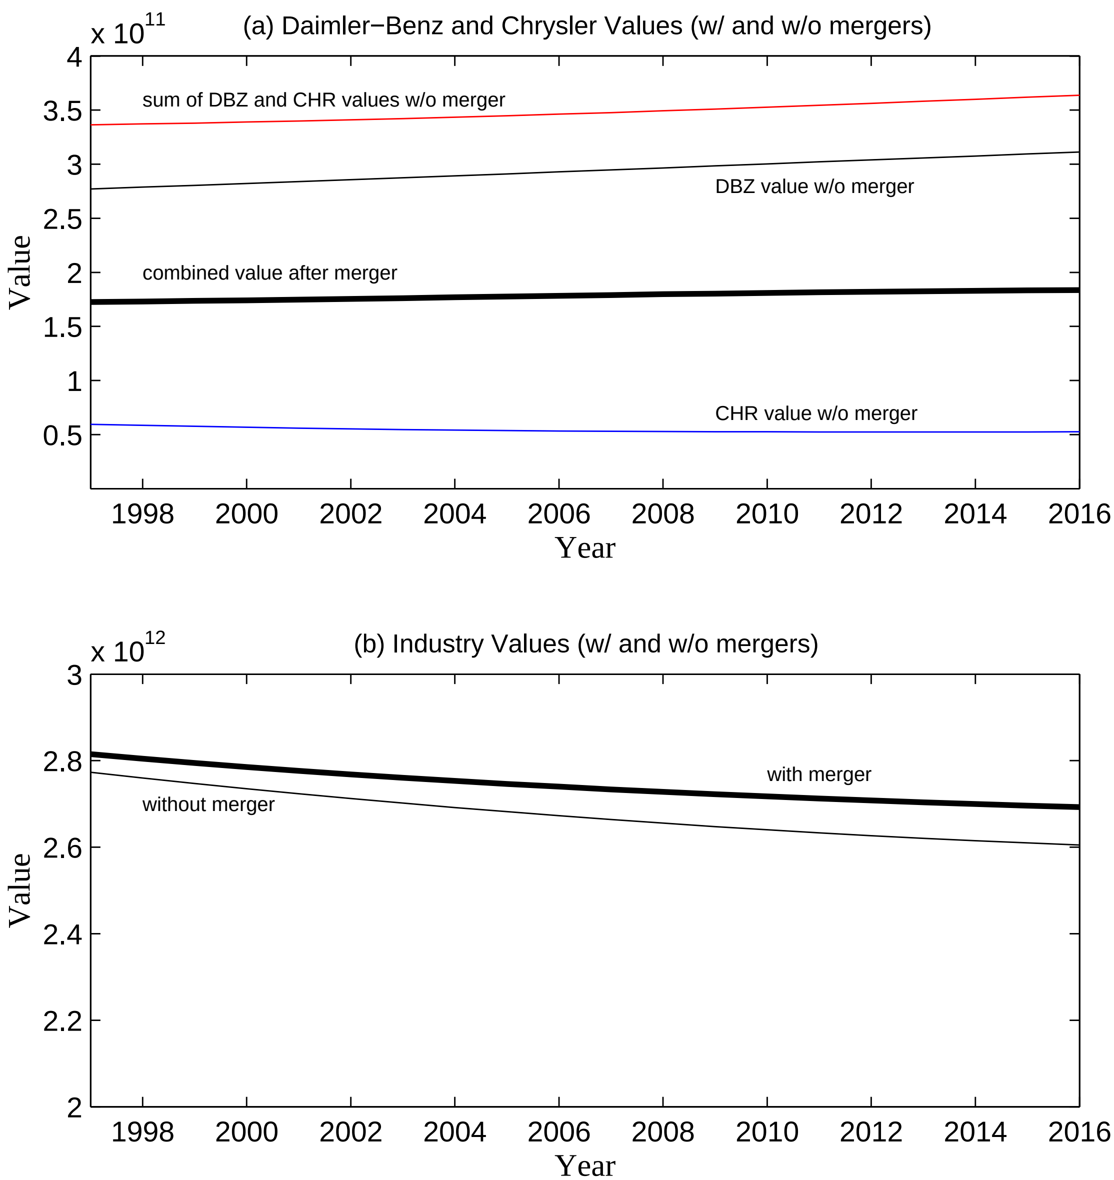

Figure 10: The Effects Daimler-Benz and Chrysler Merger on Firm and Industry Values

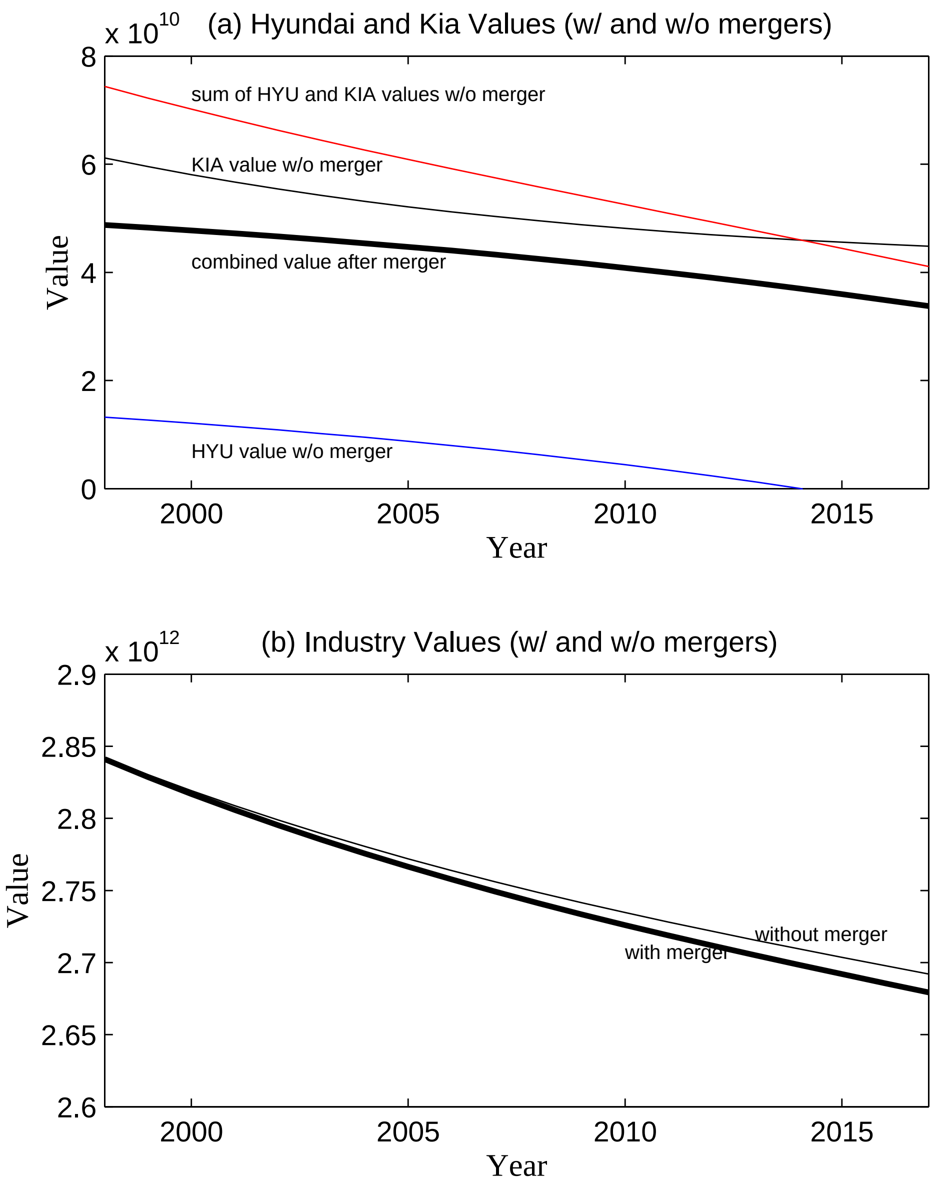

Figure 11: The Effects of Hyundai-Kia Merger on Firm and Industry Values

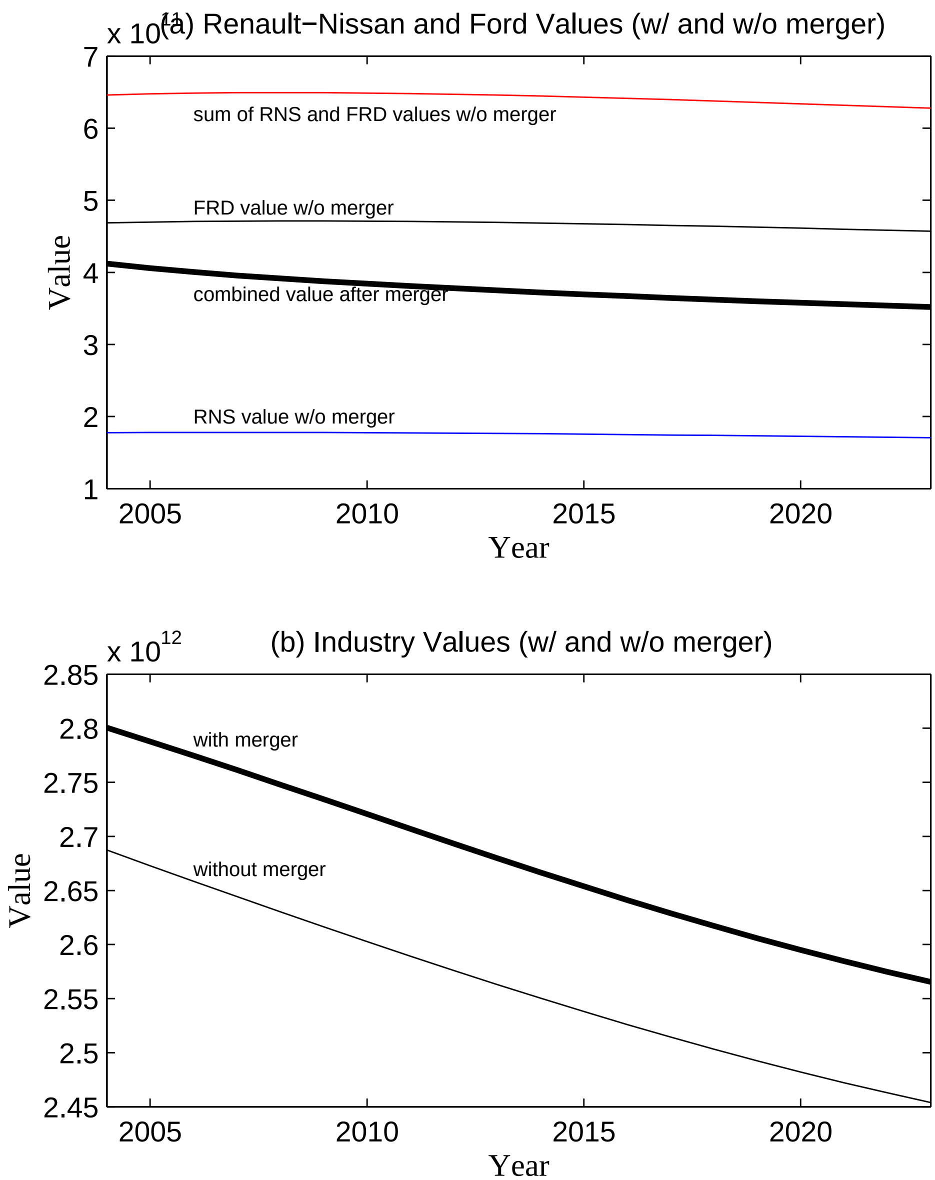

Figure 12: The Effects of Renault-Nissan and Ford Merger on Firm and Industry Values

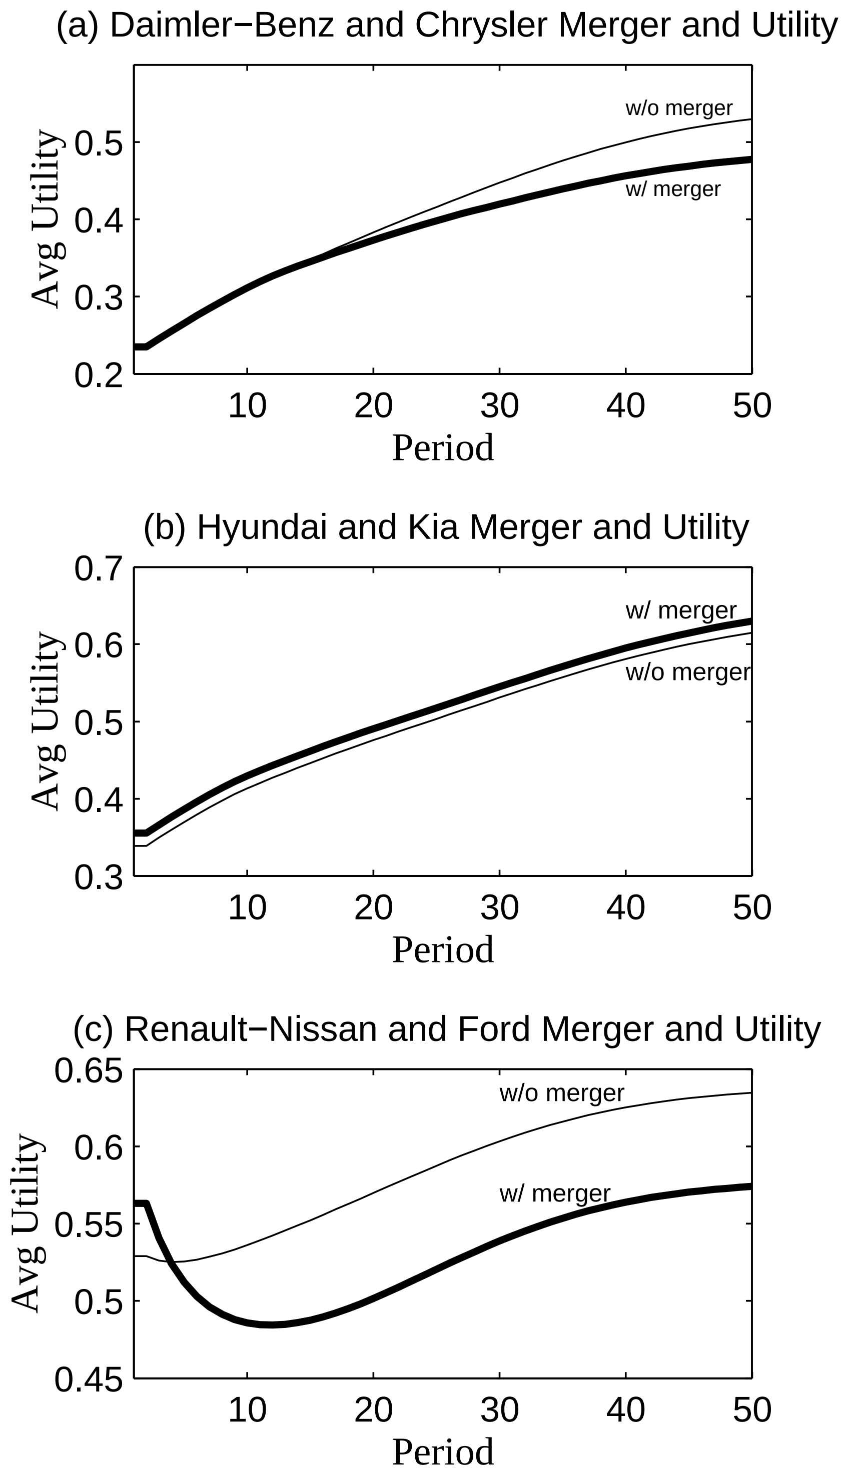

Figure 13: Mergers and Average Utility