Non-Equilibrium Industry Dynamics with Knowledge-Based Competition: An Agent-Based Computational Model

Myong-Hun Chang*
Department of Economics
Cleveland State University
Cleveland, OH 44115
216-687-4523, -9206 (Fax)
m.chang@csuohio.edu
http://academic.csuohio.edu/changm

June 2007

Abstract

An agent-based computational model is developed to explore the evolutionary dynamics of an industry which is subject to knowledge-based competition with entry and exit. It views the production process as a system of inter-dependent activities and the firm as an adaptive entity whose survival depends on its ability to perform various activities with greater efficiency than its rivals. The model is capable of generating many empirical regularities, including those on firm turnovers, evolving market structure, and the intra-industry technological diversity. Comparative dynamics analyses are performed to investigate how these regularities are affected by various industry-specific factors such as the attributes of the market environment, search propensities and the nature of the technology space in which individual firm's learning takes place. Of particular interest are the findings on firm turnover and the number of surviving incumbents in the long run. The turnover of firms is found to be greater when: 1) the market demand is larger; 2) the pool of potential entrants is larger; 3) the start-up fund for a firm is smaller; and 4) the firms in the industry have a lower propensity to search. The number of surviving incumbents is higher in the long run when: 1) market demand is larger; 2) the fixed cost is lower; 3) the pool of potential entrants is larger; and 4) the production process entails a greater number of component activities.

Keywords: Industry Dynamics, Shakeouts, Innovation, Imitation, Technological Diversity, Agent-Based Computational Model

JEL Codes: L10, O30

*I thank Jason Barr, Jon Harford, Bill Kosteas, Edward Nell, and the participants of the computational economics session at the EEA-2007 Meetings for many helpful comments. The paper has also been improved by the comments of the anonymous reviewers for the 2007 International System Dynamics Conference in Boston.

1 Introduction

The literature on industrial dynamics contains a wide array of empirical works identifying a set of regularities which arise in many manufacturing industries.[1] In particular, Gort and Klepper (1982), tracing the market histories of 46 new products, empirically verified a distinct sequence of stages that is widely observed in the development of various industries from their birth to maturity. The historical paths of the industries examined displayed a common pattern, in which the number of producers initially rose, then declined sharply, eventually converging to a stable level.[2] Examining the data on output for a subset of these industries over time, they found that the output generally grew, but the rate of growth in output declined steadily over the course of the industrial development. They also found that the market price of these products declined monotonically over time, with the greatest percentage decrease occurring during the early stage of the development. These findings have been confirmed and further elaborated upon by Klepper and Simons (1997, 2000a, 2000b), and Klepper (2002).

A small number of papers have offered highly stylized theoretical models which provide insights into the mechanisms generating the above-mentioned regularities [Klepper and Graddy (1990), Jovanovic and MacDonald (1994)]. These models entail profit-maximizing firms with foresight. In this paper, I propose an alternative model which entails a population of myopic, though adaptive, firms engaged in perpetual search for improvements in their production technologies. The approach is to develop an agent-based computational model of an industry, in which the entry/exit dynamics of the industry are driven by firm-level learning taking place in a structured search space for production technology and the market-wide competition that acts on the learning outcomes. In particular, the competition takes place on the basis of differential production efficiencies resulting from the heterogeneous stocks of knowledge (production know-hows) accumulated by the individual firms through repeated innovation and imitation attempts. The modelling of production technologies in this paper makes the representation of knowledge and its accumulation explicit and tractable.

The proposed model is capable of generating all of the empirical regularities mentioned above for a wide range of parameter configurations. The main objective of the paper is, however, to go beyond merely generating these regularities.[3] Utilizing the rich structure of the model, I engage in the comparative dynamics analyses, in which the impacts of various parameters on the resulting industry dynamics are thoroughly examined. The agent-based computational approach taken in this paper permits the type of analyses that the purely theoretical models do not. The ultimate goal is to pursue objective and coherent comparisons between industries which differ in their characteristics along the lines suggested by the parameter configurations considered in this paper. Indeed, several researchers have emphasized the need for identifying the industry-specific factors that influence such evolutionary processes:

> "A last observation concerns the enormous variation across new industries in the pace and severity of the prototypical pattern of industry evolution. This suggests

[1] See Sutton (1997) and Caves (1998) for excellent surveys.

[2] See Carroll and Hannan (2000) for additional empirical evidences.

[3] I view the ability to generate outputs having all of the empirical regularities mainly as *validating* the computational model proposed in the paper. Such model validation improves the level of confidence we have on the comparative dynamics analyses thus performed and any empirical predictions they may offer in terms of how various industry-specific factors affect the time path of the industry development.

that there are important differences across industries in the factors that condition the evolutionary process. More fundamentally, it suggests that there are exogenous factors that differ across industries that affect the pace and severity of the evolutionary process." [Klepper and Graddy (1990), p. 37]

"... we find substantial and persistent differences in entry and exit rates across industries. Entry and exit rates at a point in time are also highly correlated across industries so that industries with higher than average entry rates tend to also have higher than average exit rates. Together these suggest that industry-specific factors play an important role in determining entry and exit patterns." [Dunne, Roberts, and Samuelson (1988), p.496]

The comparative dynamics analyses in this paper are performed to respond to these calls. The industry-specific factors considered in this paper include the attributes of the market environment (such as the size of the market demand, level of the fixed cost, the availability of potential entrants, initial wealth levels of the firms), the industry-specific search propensity, and, finally, the nature of the technology space within which individual firm's learning takes place.

The existing literature on industrial dynamics already contains works with stylized theoretical models which are capable of generating the earlier-mentioned regularities. A case in point is Klepper and Graddy (1990). Their model assumes a fixed number of potential entrants, who differ in terms of their costs and product qualities. These potential entrants are profit-maximizers who enter the industry only if the expected discounted profits from entry are non-negative. Only a fraction of these potential entrants then actually enter the industry in any given period, based on their relative cost and quality positions. The incumbent firms are also heterogeneous in their costs and product qualities, where these differentials tend to persist because of the imperfection in the imitation activities of the firms. In particular, they assume that the firms may reduce their costs only through the imitation of more efficient competitors and this imitation occurs only once – immediately upon their entry into the industry. After the end of the entry period, no further cost reduction takes place. All firms are price-takers and expect the market price to decline over time. In this setting, the empirical regularity concerning the shakeout is generated through the randomness in the firms' cost draws, the improvements of the cost positions through imitation upon entry, and the eventual exits of those whose improved cost positions are not low enough to remain below the falling market price. Jovanovic and MacDonald (1994) is another work that constructs a theoretical model which is capable of generating the life cycle regularity identified by Gort and Klepper (1982). Using a stylized analytical model, where one major invention is followed by a one-time refinement of the technology, they generate the regularities concerning the shakeout process. The model is estimated using data from the U.S. automobile tire industry.

To the extent that both Klepper and Graddy (1990) and Jovanovic and MacDonald (1994) treat firm-level learning and market competition as the driving forces behind the evolution of industry structure, their works take the same perspective as the one taken in this paper. However, there are two aspects of these models that are not fully satisfactory. One, the purely theoretical approach taken in these works tend to restrict the scope of their models and analyses – e.g., the one-time-only imitation following entry in Klepper and Graddy (1990) and the one-time refinement of the technology in Jovanovic and MacDonald (1994). Two, learning in these models is represented purely in terms of its outcome – a reduction in cost – without specifying the content of the knowledge gained. Any diversity in production methods present in the population is implied solely by the existing differences in the costs. But what if the firms use two different methods of production which are equally effective? These models are not capable of differentiating between these distinct, but equally efficient, technologies. In a static setting with no learning, this is not an issue since market competition is affected purely by the firms' current cost levels. In a dynamic environment, however, in which firms are learning and accumulating their knowledge over time, the firms having heterogeneous technologies with equal efficiency levels are very likely to evolve in the long run toward those technologies having unequal efficiencies. *The exact contents of the firms' current knowledge matter in these cases, because they determine the paths the firms take in accumulating their knowledge (through innovation and imitation) and, ultimately, the evolution of the industry.* A model which treats learning purely in terms of the cost reduction, while bypassing the exact content of knowledge, will be deficient in this regard.

This paper intends to make a unique contribution to the literature by adopting a modelling approach which is capable of overcoming these limitations. There are, in fact, two methodological aspects of the chosen modelling approach which are relevant in this regard. First, the modelling approach taken here is an agent-based computational one.[4] It entails constructing a population of myopic – but adaptive – firms, each of which engages in a learning process using simple decision rules that drive (through innovation and imitation) the adoption of methods used to perform a system of production activities. These learning rules are combined with the rules governing the entry and exit behaviors of the potential entrants and the incumbent firms, respectively. By assuming a set of simple decision rules for the individual firms, we make minimal demands on the level of sophistication in their reasoning abilities.[5] Instead, the observed phenomena at the industry level are viewed as the direct consequences of the structured interactions among those decision rules which take place through market competition. All available computational resources are then dedicated to tracking such interactions and computing the time paths of the endogenous variables which characterize the behaviors of the firms and the industry. The observed regularities are treated as the realizations of these time paths for various parameter configurations.[6] It should be noted that my interest in this paper is *not* in characterizing the steady-state equilibrium, but rather the focus is on understanding the impacts of the industry-specific factors on the very process of short- to medium-run adjustments which occur on the way to such steady-states.

The second notable aspect of the model is that the decision rules involving innovation and imitation are executed in an environment which is constructed by the modeler *ex ante*.

[4] See Tesfatsion and Judd (2006) for the comprehensive and up-to-date reviews of the agent-based computational models in various economics research. The chapter by Dawid (2006) on models of innovation and technological change is particularly relevant.

[5] This approach is in line with the suggestion of Caves (1998, p. 1956) that "the evidence on entrant's growth and failure rates clearly suggests a stochastic process in which firms make their entry investments unsure of their success and do not initially position themselves at a unique optimal size."

[6] In taking this "generative" approach to explaining the macro-level phenomena, I am in perfect agreement with Epstein (2006) when he states: "... to the generativist, it does *not* suffice to demonstrate that, if a society of rational (*homo economicus*) agents were placed in the pattern, no individual would unilaterally depart – the Nash equilibrium condition. Rather, to explain a pattern, one must show how a population of cognitively plausible agents, interacting under plausible rules, could actually arrive at the pattern on time scales of interest."

This allows me to investigate the relationship between the exact structure of the environment (defined by the nature of the production technologies) and the emergent pattern of inter-firm interactions within it. More specifically, the process of innovation and imitation is viewed as that of adaptive search carried out by the boundedly rational firms in this fixed environment. The initial construction of the search environment is based on the conceptual framework which views the process of producing a good as being decomposable into a system of activities. There is a fixed number of methods which is assumed to be available for performing each activity. The choice of a production technology is defined by a vector of methods, one for each activity in the system. This vector then represents what the firm knows. In this framework, the act of innovation or imitation is viewed as search for a vector that improves the firm's production efficiency.

The production technology in this context is *complex* when the ultimate production efficiency of the firm using the technology is determined by the way these methods are combined together. This approach allows for the possibility that certain activities in the production system are complementary to one another, leading to the coexistence of multiple local optima in technology choices. Such complementarities are often formally explained using the concept of "supermodularity" [Milgrom and Roberts (1990, 1995) and Milgrom, Qian, and Roberts (1991)].[7] However, most of the theoretical work done on this concept has focused on formally defining the conditions under which such supermodularities arise.[8] They do not explore the overall impact that the degree of complementarity has on the firm's long-term search process or the resulting industry structure. This is an important omission since the degree of complementarity, which defines the complexity of the technological environment within which firms carry out their innovation/imitation activities, is what ultimately imposes the cognitive challenge on the firms, thereby affecting their abilities to compete and survive in the marketplace. The model presented in this paper is developed precisely to address this issue in a systematic fashion. The approach employed here for incorporating the inter-activity complementarity in production systems is the NK-model borrowed from Kauffman (1993), who originally used it to explore the evolutionary processes in biological systems. I utilize the NK-model in constructing the complex technological environment within which the firms search for ways to improve their production efficiencies.[9]

In the proposed model, different firms can have different sets of methods for the activities that form the system. The diversity in production methods can result in unequal production efficiencies which, in turn, lead to asymmetric average variable costs for the firms competing in the market. Such cost asymmetry generates unequal performances. When there exists a positive fixed cost of production, those firms which stay relatively inefficient tend to drop out of the market over time. This selection process is carried out while the operating firms engage in continual innovation and imitation to improve their production methods and lower

[7]Empirical support for the existence of complementarities are provided in the studies of human resource systems as exemplified in MacDuffie and Krafcik (1992), MacDuffie (1995), Ichniowski, Shaw, and Prennushi (1997), as well as a detailed case study in Ghemawat (1995). The primary purpose of this literature is to collect empirical evidences of complementarities and to identify which activities of a firm are interdependent with one another.

[8]See Vives (2005) for a comprehensive and up-to-date survey of this literature.

[9]It is worthwhile to note that the NK-model has recently been used with much success by the management scholars in exploring various organizational learning issues. See, for example, Levinthal (1997), Rivkin (2000), Rivkin and Siggelkow (2003), and Ethiraj and Levinthal (2004).

their costs. The firms in the model have limited cognitive capacity in terms of innovation and imitation, which makes it difficult for any given technology to be transferred between the firms in its entirety. When combined with the complexity in the nature of production process, this implies that there is an inherent tendency for technological diversity to persist over time. In each period, there are also potential entrants who contemplate entering the market with their own ideas on how to produce the good. Whether or not a potential entrant actually enters the market depends on how good their ideas are relative to the ones used by the existing incumbents. As the entry process takes place over time, market competition continues to apply its selection pressure on the operating firms, thereby forcing the market structure and the distribution of technologies to evolve.

The next section describes the model in detail. This is followed by a discussion in Section 3 of how the computational experiments are designed and executed. The shakeout dynamics, evolving market structure, and the intra-industry technological diversity for the baseline parameter configuration are presented in Section 4. Section 5 then explores the impacts that various classes of parameters (industry-specific factors) have on the industry dynamics identified in Section 4. Finally, a brief summary of the results and the concluding remarks are provided in Section 6.

2 The Model

The conceptual framework surrounding the proposed model is the view that a firm's production efficiency – a core determinant of its ability to compete in the market – is realized from the way its various production activities fit together as a system. Porter (1996) offers a useful example which demonstrates the competitive advantage such complementarity confers on a firm:

> Southwest's rapid gate turnaround, which allows frequent departures and greater use of aircraft, is essential to its high-convenience, low-cost positioning. But how does Southwest achieve it? Part of the answer lies in the company's well-paid gate and ground crews, whose productivity in turn-arounds is enhanced by flexible union rules. But the bigger part of the answer lies in how Southwest performs other activities. With no meals, no seat assignment, and no interline baggage transfers, Southwest avoids having to perform activities that slow down other airlines. It selects airports and routes to avoid congestion that introduces delays. Southwest's strict limits on the type and length of routes make standardized aircraft possible: every aircraft Southwest turns is a Boeing 737.... What is Southwest's core competence? Its key success factor? The correct answer is that everything matters. *Southwest's strategy involves a whole system of activities, not a collection of parts. Its competitive advantage comes from the way its activities fit and reinforce one another.* [Emphasis added.]

The model used in this paper formalizes this concept of complementarity, while viewing the process of production as a system of activities. The details of this approach are described next.

2.1 Production Process as a Complex System of Activities

A production process is composed of N distinct activities, where, for each activity, there exists a finite set of methods which can be used. For simplicity, we assume that there are exactly two methods which can be used to perform each activity. Let us represent the two methods as 0 and 1. The space of all possible production technologies is then $X \equiv \{0,1\}^N$ and a particular choice of *technology* is a binary vector of length N such that $\overline{x} \equiv (x_1, ..., x_N)$, where $x_i \in \{0,1\} \forall i$. The distance between two such vectors, $\overline{x}$ and $\overline{y}$, of length N is captured by the Hamming distance:

$$D(\overline{x}, \overline{y}) = \sum_{i=1}^{N} |x_i - y_i| \,, \tag{1}$$

which is the number of dimensions for which the vectors differ.

Associated with each technology is a numeric representing its production efficiency, $e(\overline{x})$, which is a simple average of the efficiency contributions that those N individual activities make. The crucial part of the model is how the production efficiency of a given technology is influenced by the exact way in which the methods chosen for various activities fit together. In order to address the issue of fit (or complementarity) among activities in our framework, we assume that for each activity there are $K(< N)$ other activities which influence the contribution of a given activity to the overall efficiency of the firm's production system. Let $v_i(x_i, x_i^1, ..., x_i^K)$ denote the contribution of activity i to a firm's production efficiency, where its dependence on own activity, x_i, and the K other activities to which it is coupled, $(x_i^1, ..., x_i^K)$, is made explicit.[10] For expositional convenience, we will denote by $\overline{z}_i$ the vector of methods used in the activities coupled to activity i (including itself) such that $\overline{z}_i \equiv (x_i, x_i^1, ..., x_i^K)$. We assume that the value attached to v_i for each possible vector $\overline{z}_i$ (of length $K+1$) is a random draw from $[0, 100]$ according to a uniform distribution. The overall efficiency level of a firm, when it uses $\overline{x}$, is then

$$e(\overline{x}) = \frac{1}{N} \sum_{k=1}^{N} v_k(\overline{z}_k). \tag{2}$$

Division by N is done for the purpose of normalization which facilitates comparison of the outcomes across different values of N. Clearly, $e(\overline{x}) \in [0, 100]$.

Given $e(\overline{x})$ defined for all $\overline{x} \in X$, a firm's innovation/imitation activity in our context can be viewed as *search for more efficient technology*. It is useful to think of this search as taking place on a *landscape*, which is defined on Euclidean space with each activity of a firm being represented by a dimension of the space and the final dimension indicating the efficiency of the firm. A critical factor in the firm's search process is the exact shape of this landscape. Kauffman (1993) has shown that the way various activities interact with one another through N and K determines the shape of the landscape. The main result is that for $K > 0$ the landscape is rugged so that there are multiple local optima. Furthermore, the average number of local optima on such rugged landscapes tends to increase in N and K – the technological environment is more complex for higher values of N and/or K. These properties will become crucial later on when we address the issue of technological diversity in an evolving industry. It

[10] In the computational experiments, the K activities to which activity i is coupled are chosen randomly from $(N-1)$ other activities with a uniform distribution.

is, hence, through varying the values of N and K that we can directly control the nature of the technological environment within which the boundedly rational firms carry out their search.

Since the methods vectors evolve over time as firms successfully engage in search, we put the time superscript on them for expositional clarity: $\overline{x}_i^t \equiv (x_{i,1}^t, x_{i,2}^t, ..., x_{i,N}^t)$ denotes the production technology employed by firm i in period t, where $\overline{x}_i^t \in X$ and $x_{i,k}^t \in \{0,1\}$ is the method chosen for activity k. Firm i's production efficiency level in period t is then $e(\overline{x}_i^t)$.

2.2 Demand and Cost

In each period, there exists a finite number of firms that operate in the market. In this section, we define the static market equilibrium among the operating firms. As such, we will abstract away from the time superscript temporarily. Let m be the number of firms in the market. The market is that of a homogeneous good. The firms are Cournot oligopolists, where they choose production quantities: q_i denotes firm i's production quantity. In defining the Cournot equilibrium in this setting, we assume that all m firms produce positive quantities in equilibrium.[11] The inverse market demand function is specified to be

$$P(Q) = a - Q, \tag{3}$$

where $P(Q)$ is the price consumers are willing to pay when there are $Q(=\sum_{j=1}^m q_j)$ units of total output available in the market.

Each operating firm has its production technology, $\overline{x}_i$, and faces the following total cost function:

$$C(q_i) = f_i + c_i \cdot q_i. \tag{4}$$

Hence, f_i is a fixed cost of production for firm i, while c_i is its marginal cost. The firm's marginal cost depends on the level of production efficiency embedded in the technology that the firm is using. More specifically, I assume that c_i is a declining function of $e(\overline{x}_i)$ and specify the following simple form

$$c_i(\overline{x}_i) = 100 - e(\overline{x}_i). \tag{5}$$

The total cost can be re-written as:

$$C(q_i) = f_i + (100 - e(\overline{x}_i)) \cdot q_i. \tag{6}$$

2.3 Cournot Equilibrium with Asymmetric Costs

Given the inverse market demand function and the firm cost function, firm i's profit is:

$$\pi_i(q_i, Q - q_i) = \left(a - \sum_{j=1}^{m} q_j\right) \cdot q_i - f_i - c_i \cdot q_i \tag{7}$$

For simplicity, let us assume that the firms have identical fixed cost so that $f_1 = f_2 = ...f_m \equiv f$.

[11] In actuality, there is no reason to suppose that in the presence of asymmetric costs all m firms will produce positive quantities in equilibrium. Some of these firms may become *inactive* by producing zero quantity. The algorithm used to distinguish among active and inactive firms based on their production costs will be addressed in a later section.

The first-order conditions for profit maximization imply that

$$\left(a - \sum_{j=1}^{m} \widehat{q}_j\right) - \widehat{q}_i - c_i = 0 \tag{8}$$

$\forall i \in \{1, ..., m\}$, where the output vector in Cournot equilibrium is $(\widehat{q}_1, \widehat{q}_2, ..., \widehat{q}_m)$. Adding the first-order conditions for all firms yields

$$m \cdot \left(a - \sum_{j=1}^{m} \widehat{q}_j\right) - \sum_{j=1}^{m} \widehat{q}_j = \sum_{j=1}^{m} c_j \tag{9}$$

Dividing both sides by m and simplifying, we get

$$\sum_{j=1}^{m} \widehat{q}_j = a\left(\frac{m}{m+1}\right) - \left(\frac{1}{m+1}\right)\sum_{j=1}^{m} c_j \tag{10}$$

Hence, the equilibrium market output (and the equilibrium market price) depends only on the *sum* of the marginal costs and not on the *distribution* of c_is [Bergstrom and Varian (1985)]. Using the inverse demand function, one can then write the equilibrium market price as $\widehat{P}$, where

$$\widehat{P} = \left(\frac{1}{m+1}\right)\left(a + \sum_{j=1}^{m} c_j\right) \tag{11}$$

Given the vector of marginal costs defined by the firm's chosen technology, $\widehat{P}$ is then uniquely determined. Furthermore, from the first order condition for each firm one can then express the Cournot equilibrium output rate as

$$\widehat{q}_i = \widehat{P} - c_i \tag{12}$$

$$= \left(\frac{1}{m+1}\right)\left(a + \sum_{j=1}^{m} c_j\right) - c_i \tag{13}$$

$\forall i \in \{1, ..., m\}$. A firm's equilibrium output rate depends on its own marginal cost and the sum of all marginal costs. The Cournot equilibrium firm profit is

$$\pi(\widehat{q}_i) = \widehat{P}\widehat{q}_i - f - c_i\widehat{q}_i \tag{14}$$

$$= \left(\widehat{P} - c_i\right)\widehat{q}_i - f$$

$$= (\widehat{q}_i)^2 - f \tag{15}$$

Note that $\widehat{q}_i$ is a function of c_i and $\sum_{j=1}^{m} c_j$, while c_k is a function of $\overline{x}_k$ for all k. It is then straightforward that the equilibrium firm profit is fully determined, once the vectors of methods for all firms are known. Further note that $c_i \leq c_k$ implies $\widehat{q}_i \geq \widehat{q}_k$ and, hence, $\pi(\widehat{q}_i) \geq \pi(\widehat{q}_k) \forall i, k \in \{1, ..., m\}$.

2.4 Dynamic Structure of the Model

In every period of the horizon in this model, there are four stages of firm decision making. Figure 1 shows the sequence of these decision stages. The definitions of the set notations introduced in this section and used throughout the paper are summarized in Table 1.

The process of intra-industry dynamics in period t starts with four groups of state variables. First, there exists a set of surviving firms from $t-1$, denoted S^{t-1}, where $S^0 = \emptyset$. The set of surviving firms includes those firms which were *active* in $t-1$ in that their outputs were strictly positive as well as those firms which were *inactive* with their plants shut down during the previous period. Let S_a^{t-1} and S_{-a}^{t-1} denote, respectively, the set of *active* and *inactive* firms in $t-1$ such that

$$S_a^{t-1} \equiv \{\text{all } j \in S^{t-1} | q_j^{t-1} > 0\}, \tag{16}$$

$$S_{-a}^{t-1} \equiv \{\text{all } j \in S^{t-1} | q_j^{t-1} = 0\}, \tag{17}$$

and $S_a^{t-1} \cup S_{-a}^{t-1} = S^{t-1}$. The inactive firms in $t-1$ are able to survive if they have sufficient wealth balance to cover their fixed costs for that period.

Second, each firm $i \in S^{t-1}$ possesses a production technology, $\overline{x}_i^{t-1}$, carried over from $t-1$, which gives rise to its current efficiency level of $e(\overline{x}_i^{t-1})$ and the consequent marginal cost of c_i^{t-1} as defined in equation (5).

Third, each firm $i \in S^{t-1}$ has a current wealth level of w_i^{t-1} which it carries over from $t-1$. This wealth level is adjusted at the end of each period on the basis of the economic profit earned (which adds to it) or loss incurred (which subtracts from it) by the firm. It is this wealth level which ultimately determines the firm's viability in the market.

Finally, there is a finite set of *potential* entrants, R^t, who contemplate entering the industry in the beginning of t. In this paper, we assume that the size of the potential entrants pool is fixed and constant at r throughout the entire horizon. We also assume that this pool of r potential entrants is renewed fresh each period. Each potential entrant k in R^t is endowed with a technology, $\overline{x}_k^t$, randomly chosen from X according to uniform distribution. Associated with the technology is its corresponding efficiency level, $e(\overline{x}_k^t)$, and the marginal cost of c_k^t for all $k \in R^t$.

2.4.1 Stage 1: Entry Decisions

In stage 1 of each period, the potential entrants in R^t first make their decisions to enter. This entry decision is dependent upon the common threshold efficiency level, $\widehat{e}^t$, which is defined as follows:

$$\widehat{e}^t = \begin{cases} 0 & \text{for } t = 1, \\ \min\{e(\overline{x}_i^{t-1})\}_{i \in S_a^{t-1}} & \text{for } t > 1. \end{cases} \tag{18}$$

$\widehat{e}^t$ is, hence, the efficiency level of the least efficient *active* incumbent from the previous period, except for $t = 1$ when there is no past record of efficiency levels to compare to. Given the threshold level, the decision rule of a potential entrant $k \in R^t$ is:

$$\begin{cases} \text{Enter,} & \text{if and only if } e(\overline{x}_k^t) \geq \widehat{e}^t \\ \text{Do not enter,} & \text{otherwise.} \end{cases} \tag{19}$$

The decision rule indicates that an outsider will be attracted to enter the industry if and only if it is convinced that it is at least as efficient as the least efficient incumbent who *actively* engaged in production in the previous period. Implicit behind this assumption is that the efficiency level of an incumbent firm can be inferred by an outsider only if it produces a positive output. Note from equation (12) that the least efficient active incumbent firm is the one producing the minimum positive output. Assuming that the potential entrants have the knowledge of the cost structure as specified in equation (5), they only need to observe the outputs and the market price to correctly infer the efficiency levels of the incumbent firms.[12]

Once every potential entrant in R^t makes its entry decision on the basis of the above criterion, the resulting set of *actual* entrants, $E^t \subseteq R^t$, contains only those firms with sufficiently efficient technologies. Denote by M^t the set of firms ready to compete in the industry: $M^t = S^{t-1} \cup E^t$. We will denote by m^t the number of competing firms in period t such that $m^t = \left|M^t\right|$.

All entrants into the industry enter with a fixed "start-up budget" of b dollars. The start-up budget may be viewed as a firm's available fund that remains after paying for the one-time setup cost of entry.[13] For example, if one wishes to consider a case where a firm has zero fund available, but must incur a positive entry cost, it would be natural to consider b as having a strictly negative value. Just as each firm in S^{t-1} has its current wealth level of w_i^{t-1}, we will let $w_j^{t-1} = b$ for all $j \in E^t$. At the end of stage 1 of period t, we then have a well-defined set of competing firms, M^t, and the current wealth levels for all firms in that set, $\{w_i^{t-1}\}_{\forall i \in M^t}$.

2.4.2 Stage 2: Innovation/Imitation Decisions

In stage 2, the surviving incumbents from $t-1$ – i.e., all firms in S^{t-1} – engage in innovation or imitation in order to improve the efficiency of their existing technologies. Given that the new entrants in E^t, selected in stage 1, entered with new technologies, they do not engage in the innovation/imitation process in t.

Each surviving incumbent in t gets one chance to *search* (through innovation or imitation) in that period with probability α. With probability $(1-\alpha)$, it does not get the opportunity to search, in which case $\overline{x}_i^t = \overline{x}_i^{t-1}$. With α, we are then capturing the exogenously specified firm propensity (common to all firms) to engage in exploration for improved technologies in a given industry.[14] This opportunity (incentive) may be determined by the institutions exogenous to the market competition or by the prevailing culture of innovation within the industry that is not specified in the model.

If firm i gets to search in period t, it then chooses to "innovate" with probability β_i^t and "imitate" with probability $1-\beta_i^t$. [The probability β_i^t is endogenous – how it is updated from one period to the next is discussed below.] Innovation occurs when the firm considers changing

[12] Two comments are in order here. One, the potential entrants can infer the overall efficiency levels of the active incumbents, but not the individual components of the efficiency vector. Two, a potential entrant correctly infers its own efficiency level given the technology it is endowed with. While a more general model may specify a noisy evaluation of the efficiency level, we avoid complicating the analysis by simply assuming that the potential entrants are already familiar with their technologies. This is likely if they happen to be the established players in other markets, trying to expand into the market in question.

[13] The size of the one-time cost of entry is not directly relevant for our analysis. It may be zero or positive. If it is zero, then b is the excess fund the firm enters the market with. If it is positive, then b is what remains of the fund after paying for the cost of entry.

[14] In reality, the search propensity, α, is likely to be heterogeneous among firms. It may also be endogenous in that it may evolve over time on the basis of individual firm's realized profits. I specify it to be homogeneous and exogenous in this model so as to keep our computational analysis tractable.

the method in *one* randomly chosen activity. Imitation occurs when the firm picks another firm j from a subset of S^{t-1} and considers copying the method employed by j in *one* randomly chosen activity.[15] Only those surviving firms which were profitable in $t-1$, i.e., $\pi_k^{t-1} > 0$, are considered as the potential targets for imitation. Let S_b^{t-1} denote the set of these *profitable* firms, where $S_b^{t-1} \subseteq S^{t-1}$. The choice of a firm to imitate is made probabilistically using the "roulette wheel" algorithm. To be specific, the probability of firm $i \in S^{t-1}$ observing a firm $j \in S_b^{t-1}$ is denoted p_{ij}^t and is defined as follows:

$$p_{ij}^t \equiv \frac{\pi_j^{t-1}}{\sum_{\forall k \in S_b^{t-1}, k \neq i} \pi_k^{t-1}},\ j \neq i \tag{20}$$

such that $\sum_{\forall j \in S_b^{t-1}, j \neq i} p_{ij}^t = 1 \forall i \in S^{t-1}$. Hence, the more profitable firm is more likely to be observed and imitated.

Let $\widetilde{x}_k^t$ denote firm k's vector of experimental methods (i.e., a technology considered for potential adoption) obtained through "innovation" or through "imitation." The adoption decision rule is as follows:

$$\overline{x}_k^t = \begin{cases} \widetilde{x}_k^t, & \text{if and only if } e(\widetilde{x}_k^t) > e(\overline{x}_k^{t-1}) \\ \overline{x}_k^{t-1}, & \text{otherwise.} \end{cases} \tag{21}$$

A proposed technology is adopted by a firm if and only if the resulting efficiency level exceeds that of its current technology. Implicit behind this rule is the assumption that the firm is able to correctly infer the efficiency level associated with the experimental technology, $\widetilde{x}_k^t$. While we believe that the firm can obtain such information only through actual experiments, we abstract away from modelling such experiments so as not to overload the already complex model.

We intentionally limit the scope of technological change to one activity at a time. This implies a cognitive constraint faced by the firms. As the number of activities in which methods can be changed rises above one, the firms are then able to consider innovations and imitations of larger scale and the existence of complementarities among activities (and the resulting multiplicity of local optima) poses no problem in the long run since the firms are able to jump from one local optimum to another.[16]

Let us now get back to the choice probability, β_i^t. In our setting, α (search propensity) is exogenous and common to all firms, while β_i^t is endogenous and specific to each firm. More specifically, the choice probabilities of β_i^t and $1-\beta_i^t$ are adjusted over time by individual firms according to a reinforcement learning rule. We adopt a version of the *Experience-Weighted Attraction (EWA)* learning rule as described in Camerer and Ho (1990). Under this rule, a firm has a numerical attraction for each possible action – *innovation* or *imitation* in our case. The learning rule specifies how attractions are updated by the firm's experience and how the probabilities of choosing different actions depend on attractions. The main feature of the rule

[15] Hence, the imitating firm is assumed to be capable of copying only a small part of the entire technology.

[16] An additional constraint in terms of a firm's cognitive capacity is the degree of precision in its evaluation of the production efficiency. So as to avoid overloading the model, we assume that the production efficiency of a firm's technology is assessed with perfect accuracy.

is that a positive outcome realized from a course of action reinforces the likelihood of that same action being chosen again.

Using the *EWA*-rule, β_i^t is adjusted each period on the basis of evolving attraction measures, $B_i^{IN}(t)$ for innovation and $B_i^{IM}(t)$ for imitation. The evolution of $B_i^{IN}(t)$ and $B_i^{IM}(t)$ follow the process below:

$$B_i^{IN}(t+1) = \begin{cases} \phi B_i^{IN}(t) + 1, & \text{if firm } i \text{ adopted a technology through innovation in } t \\ \phi B_i^{IN}(t), & \text{otherwise} \end{cases} \tag{22}$$

$$B_i^{IM}(t+1) = \begin{cases} \phi B_i^{IM}(t) + 1, & \text{if firm } i \text{ adopted a technology through imitation in } t \\ \phi B_i^{IM}(t), & \text{otherwise} \end{cases} \tag{23}$$

where $\phi \in (0,1]$ is the decay factor. Hence, if the firm chose to pursue *innovation* and discovered and then adopted a new idea, the attraction measure for innovation increases by 1 after allowing for the decay factor of ϕ on the previous attraction level. If the firm chose innovation but was unsuccessful (because the idea generated was not useful) or if it instead chose imitation, then its new attraction measure for innovation is simply the attraction level from the previous period decayed by the factor ϕ. Similarly, a success or failure in imitation at t has identical influence on $B_i^{IM}(t+1)$. For analytical simplicity, we assume $\phi = 1$ throughout this paper so that the attractions do not decay.

Given $B_i^{IN}(t)$ and $B_i^{IM}(t)$, one derives the choice probability of innovation in period t as follows:

$$\beta_i^t = \frac{B_i^{IN}(t)}{B_i^{IN}(t) + B_i^{IM}(t)}. \tag{24}$$

The probability of pursuing imitation is, of course, $1 - \beta_i^t$. The expression in (24) implies that a favorable experience through innovation raises the probability that a firm will choose innovation again in the future.

2.4.3 Stage 3: Output Decisions and Market Competition

Given the consequences of the innovation/imitation choices made in stage 2 by the firms in S^{t-1}, all firms in M^t now have the updated technologies $\{\overline{x}_i^t\}_{\forall i \in M^t}$ as well as their current wealth levels $\{w_i^{t-1}\}_{\forall i \in M^t}$. The updated technologies define the efficiency levels of the firms, $\{e(\overline{x}_i^t)\}_{\forall i \in M^t}$, and their corresponding marginal costs for period t, $\{c_i^t\}_{\forall i \in M^t}$. Given these marginal costs, the firms engage in Cournot competition in the market, where we approximate the outcome with the Cournot equilibrium defined in 2.3.

Note that the equilibrium in 2.3 was defined for m firms who were assumed to produce positive quantities in equilibrium. In actuality, given the asymmetric costs, there is no reason to think that all m^t firms will produce positive quantities in equilibrium. Some relatively inefficient firms may shut down their plants and stay inactive. What we need is then a mechanism for identifying the set of *active* firms out of M^t such that the Cournot equilibrium among these firms will indeed entail positive quantities only. This is accomplished in the following sequence of steps. Starting from the initial set of active firms, compute the equilibrium outputs for each firm. If the outputs for one or more firms are negative, then de-activate the least efficient firm from the set of currently active firms – i.e., set $q_i^t = 0$ where i is the least efficient firm. Re-define the set of active firms (as the previous set of active firms minus the de-activated firms) and recompute the equilibrium outputs. Repeat the procedure until all active firms are producing non-negative outputs. Each *inactive* firm produces zero output and incurs the economic loss equivalent to its fixed cost. Each *active* firm produces its Cournot equilibrium output and earns the corresponding profit. We then have π_i^t for all $i \in M^t$.

2.4.4 Stage 4: Exit Decisions

Given the single-period profits or losses made in stage 3 of the game, the incumbent firms consider exiting the industry in the final stage. The incumbent firms' wealth levels are first updated on the basis of the profits (losses) made in t:

$$w_i^t = w_i^{t-1} + \pi_i^t. \tag{25}$$

The exit decision rule for each firm is:

$$\begin{cases} \text{Stay in} & \text{iff } w_i^t \geq d, \\ \text{Exit} & \text{otherwise,} \end{cases} \tag{26}$$

where d is the threshold wealth level such that all firms with their current wealth levels below d exit the market. Once the exit decisions are made by all firms in M^t, the set of surviving firms from period t is then defined as:

$$S^t \equiv \{\forall i \in M^t | w_i^t \geq d\}. \tag{27}$$

We denote by L^t the set of firms which have decided to exit:

$$L^t \equiv \{\forall i \in M^t | w_i^t < d\}. \tag{28}$$

The set of surviving firms, S^t, their current technologies, $\{\overline{x}_i^t\}_{\forall i \in S^t}$, and their current wealth levels, $\{w_i^t\}_{\forall i \in S^t}$, are then passed on to $t+1$ as the state variables.

3 Design of Computational Experiments

The ultimate objective is to examine the time paths of certain endogenous variables for various parameter configurations that are relevant for the evolution of the industry. The definitions as well as the values of the parameters considered in the simulations are provided in Table 2.

The endogenous variables I am interested in can be separated into three categories. First, there are those which characterize the turnover of the firms. These include: 1) $\left|M^t\right|$, the number of all operating firms in t; 2) $\left|E^t\right|$, the number of actual entrants in t; 3) $\left|L^t\right|$, the number of exiting firms in t.

The second category of endogenous variables include those capturing the state of the market: 1) $\widehat{P}^t$, the equilibrium market price in t; 2) $\widehat{Q}^t$, the equilibrium industry output, where

$$\widehat{Q}^t \equiv \sum_{i=1}^{m^t} \widehat{q}_i^t; \tag{29}$$

3) h^t, the Herfindahl-Hirschmann Index in t, where

$$h^t \equiv \sum_{i=1}^{m^t} \left(\frac{\widehat{q}_i^t}{\widehat{Q}^t} \cdot 100 \right)^2 . \tag{30}$$

For lack of space, we will characterize the behavior of the second category variables only for the baseline case. The comparative dynamics results pertaining to these variables will not be reported here.

Finally, the third category describes the evolution of the technological diversity within the industry. The relevant outputs to examine are the distributions of the operating firms' marginal costs, outputs, and their technologies: $\{c_i^t\}_{\forall i \in M^t}$, $\{q_i^t\}_{\forall i \in M^t}$, $\{\overline{x}_i^t\}_{\forall i \in M^t}$. Since some firms may have identical technologies, identical marginal costs, and, hence, identical equilibrium outputs, one can summarize the extent of diversity among firms at any given point in time by computing the proportion of all operating firms who have distinct technologies.[17]

We run 1000 replications for each parameter configuration, using a fresh set of random numbers for each replication.[18] We report the time paths of the above variables directly from a single typical replication as well as those averaged over 1000 replications. Outputs are examined for a sufficiently large number of individual replications so as to ensure that the time paths from a single replication reported in the paper are indeed typical of all replications for a given parameter configuration.[19]

4 The Baseline Case: Generating and Identifying Regularities

The baseline case considered in this section assumes the following configuration of parameter values: $N = 16$; $K = 2$; $r = 10$; $f = 20$; $a = 200$; $b = 100$; $d = 0.0$; $\alpha = 1.0$; $T = 4,000$. The manufacturing process has 16 component activities, where the efficiency contribution of each activity is determined by the method chosen for that activity and the methods chosen for *two* other activities which are directly linked to it. There are 10 potential entrants contemplating entry into the market each period. The fixed cost is 20 and the demand intercept is 200. Each firm, when it enters the industry, has a start-up budget of 100. An existing incumbent leaves the market if its current wealth balance falls below 0.0. The operating firms engage in search every period, where the search mechanism is chosen between innovation and imitation based on the probabilities that evolve over time through reinforcement learning. The cognitive skills of the firms are such that they are capable of evaluating the efficiency consequence of changing the method in only *one* activity out of N.

Before reporting the full set of results for the baseline case, I first show that the model presented in this paper is indeed capable of generating the empirical regularities mentioned earlier. The key empirical regularities involve the time paths of 1) the number of producers in the industry, 2) the aggregate industry output and 3) the market price. Using the data on

[17]Given that a technology in this model is a vector of 0s and 1s, it is quite simple to distinguish among different technologies.

[18]Since we consider a horizon of 4,000 periods in each of our replications, we are then computing the outcomes from 4 million occurrences of the market competition for each parameter configuration.

[19]The source code for the computational experiments was written in C++ and the simulation outputs were analyzed and visualized using Mathematica 3.0. The source code is available upon request from the author.

U.S. automobile tire industry, initially collected by Gort and Klepper (1982) and published in Jovanovic and MacDonald (1994), I plot in Figure 2 the following: 1) the number of producers over 68 years from 1906 to 1973, 2) the industry output over 64 years from 1910 to 1973, and 3) the price index over 61 years from 1913 to 1973. The complete data sets are provided in the Appendix of Jovanovic and MacDonald (1994).

As initially shown by Gort and Klepper (1982), the number of producers rises sharply in the beginning, reaching the maximum of 275 in 1922. It then declines sharply, eventually leveling off to a stable level. The industry output appears to rise and the wholesale price index declines over time.

Given the empirical results displayed in Figure 2, I now compare them to the computational results generated with our model for the baseline parameter configuration. Figure 3 shows the time paths of the three variables – number of producers, industry output, and market price – from a single typical replication. These time paths are plotted only for the first 68 periods in order to facilitate comparisons with the empirical data plotted in Figure 2. The qualitative similarities between the figures in Figure 2 and Figure 3 are striking.[20] The number of incumbents in the baseline model rises sharply in the beginning, reaches a maximum, and then turns downward, eventually stabilizing to a moderate level. The industry output grows just as in the case of the automobile tire industry, but the rate of its growth declines over time. This is in contrast to the output path of the tire industry plotted in Figure 2 – the output appears to increase at a slightly increasing rate in that case. However, the simulated output path in Figure 3 is in perfect agreement with the general pattern observed in terms of the percentage change in output for 25 different products as reported in Gort and Klepper (1982) and Klepper and Graddy (1990): "The summary statistics ... suggest that the rate of growth of output declines steadily over the course of the development of new product industries." [Gort and Klepper (1982), pp. 644-645] The market price path generated by the model also has the same general shape as that created by the actual data. Both the simulated and the actual price paths appear to fall at a decreasing rate.

4.1 Entry, Exit, and Shakeout

I now present the full set of results that are generated in the baseline case for the entire horizon of 4,000 periods. We start with the processes of entry and exit and the resulting market structure which evolves over time. Figure 4(a) reports the time paths over $t \in \{1, ..., 4000\}$ of the following endogenous variables from a single typical replication: 1) number of incumbent firms, m^t, 2) number of actual entrants, $|E^t|$, and 3) number of exits, $|L^t|$. These measures were then obtained for 1,000 independent replications and averaged. The time paths of those averages are reported in Figure 4(b). Both sets of figures are drawn with the logarithm of time index along the horizontal axis so that the adjustments taking place in the first 100 periods are magnified relative to those in later periods. [In fact, the same scaling will be used for most of the figures presented in this paper.]

The most striking property observed in these figures is the one involving the number of incumbent firms, m^t. Given that the size of the potential entrants pool is fixed at 10 each period, the number of incumbent firms starts at $m^1 = 10$. Since the firms enter with a

[20] Since my interest in this paper is in examining the qualitative nature of the dynamics, I make no attempt to calibrate the model to improve the fit between the computational and empirical results.

fixed budget of $b = 100$, it takes some time before any exit occurs – the relatively inefficient firms must accumulate sufficient economic losses before they go bankrupt. In the meantime, entry continues as those potential entrants with good draws on technology find themselves sufficiently efficient to enter and compete in the industry. This is shown in the time paths of $\left|E^t\right|$ captured in Figure 4, where the number of actual entrants, even as it declines, remains rather high in the beginning. For the first 10 periods or so, the total number of incumbent firms rises because of this inflow of the new entrants. The increase in m^t, however, comes to a stop eventually. As can be seen in the bottom figures of Figure 4(a) and 4(b), the time path of the number of exits starts to rise after the fifth period. These exits of the relatively inefficient firms occur as the general reduction in costs achieved through innovation and imitation puts the downward pressure on the market price. The combination of diminishing entry and rising exit then leads to a rapid decline in the number of incumbent firms, which eventually comes to a somewhat stabilized level at around $t = 50$.

The shakeout phenomenon observed in so many new industries arises naturally in this model, in which the opening-up of a new market is followed by firms entering with operating practices of various efficiency levels. While the new available market demand invites entry, the ensuing market competition applies sufficient selection pressure on the firms of unequal production efficiencies that the exits of relatively inefficient firms become inevitable. The simulation outputs not reported here show that shakeout is a general phenomenon that arises for a wide variety of parameter configurations considered in this paper.

4.2 Output, Price, and Concentration

As mentioned earlier in this section, the movements of the market price and the industry output have been shown to exhibit certain regularities. Figure 5 captures the time paths of the market price and the industry output that are generated with the baseline parameter configurations – 5(a) from a single typical replication and 5(b) from averaging of the 1,000 independent replications. The results confirm the empirical observations made by Gort and Klepper (1982) as well as Klepper and Graddy (1990).

Based on each firm's output, $\widehat{q}_i^t$, I also computed the Herfindahl-Hirschmann Index (HHI) each period for all $t \in \{1, ..., T\}$. The resulting time path, h^t, is displayed at the bottom of Figure 5. The HHI declines over time such that the market becomes increasingly less concentrated. In particular, the HHI declines steeply during the first 10 periods, reflecting the rapid increase in m^t in this phase as captured in Figure 4. The decline slows down after the initial 10 periods and, between $t = 10$ and $t = 100$, it tends to go back up slightly or remain constant. Recall from Figure 4 that the number of incumbents, m^t, declines significantly between $t = 10$ and $t = 100$. Purely on the basis of the number of firms alone, one would expect the HHI to rise during these periods. This is, indeed, what is shown in Figure 5(a). On average, however, this is not necessarily what happens. According to Figure 5(b) in which the HHI values are averaged over 1,000 replications, the time path of the HHI between $t = 10$ and $t = 100$ stays constant, even though the number of firms in the industry drops during the same period. Why is this so? The answer lies in the asymmetry in the firms' market shares caused by the asymmetry in marginal costs. How the technological diversity giving rise to such asymmetries evolves over time is the focus of the discussion in the next section.

4.3 Technological Diversity

Note that the firms initially start out with technologies which are chosen randomly. Clearly, the initial marginal costs are asymmetric based on this fact alone. The main question, however, is whether the firms will eventually come to adopt a common technology and, hence, converge on a common level of marginal cost as they innovate (or imitate) over time in search of more efficient technologies.

Let us first look at the marginal costs of all operating firms in M^t for all $t \in \{1, ..., T\}$. Figure 6(a) captures how the distribution of marginal costs evolves over time using the data generated from a single replication. Clearly, the distribution shifts down over time as the firms successfully innovate and adopt improved operating methods. Furthermore, the range of the marginal costs in t narrows as t increases, indicating that technological convergence is indeed taking place among firms. Figure 6(b) plots the evolving distribution of firm outputs from the same replication. These are the Cournot equilibrium outputs based on the distribution of marginal costs captured in 6(a). Note that the baseline configuration of parameters used to generate Figure 6 entails $K = 1$. The technology landscape is then moderately rugged. The single simulation reported in Figure 6 shows that the firms converge to 5 different marginal cost levels at the end of the horizon.

In order to capture the intra-industry technological diversity and its evolution in the most direct manner, the total number of distinct technologies – i.e., non-identical methods vectors – held by all operating firms in each period is counted and its time path is reported in Figure 7. The time path from a single replication is reported in Figure 7(a), while the average of those from 1,000 replications is reported in Figure 7(b). Note that the number of distinct technologies starts out at 10 in $t = 1$, which implies that all initial entrants tend to start out with distinct technologies. As more firms enter, there is a proliferation of distinct technologies between $t = 1$ and $t = 10$. With the decline in the number of firm following the peak at around $t = 10$, the number of distinct technologies also declines.

It is important to note here that the number of distinct technologies in a given period tends to fall strictly below the number of incumbent firms [compare Figure 4(a) and Figure 7(a) – the number of incumbent firms stabilizes to around 27, while the number of distinct technologies declines to 5]. This implies that there is indeed a substantial degree of technological convergence occurring over time as firms innovate (or imitate) and move toward some common local optima in the search space. The convergence is, of course, not perfect when $K > 0$. In the single replication captured in Figure 7(a), there are five distinct technologies in existence at $t = 4,000$. As shown in Figure 7(b), the number of distinct technologies at $t = 4,000$ is, on average, about 4. It tends to have a rather wide distribution across replications. Figure 8 shows the histogram of the number of distinct technologies at $t = 4,000$ over 1,000 independent replications. They range anywhere from 1 to 17, although in over 40% of the replications there are only one or two distinct technologies. Diversity among firms is then not just a transient phenomenon, but is more an inherent feature of the industries that are being modelled here.

The fact that the time paths in Figure 4(a) and Figure 7(a) have similar shapes implies that the number of distinct technologies depends on the number of operating firms, which changes over time. Hence, the initial sharp rise in the number of firms will automatically give rise to a proliferation of distinct technologies. Since we look for a measure of technological diversity which is consistent over time and is capable of controlling for the intertemporal changes in the number of firms, we define a measure that takes the ratio of the number of distinct technologies to the number of operating firms in period t:

$$\delta(t) = \frac{\text{number of distinct technologies in } t}{\text{number of operating firms in } t\ (\equiv m^t)} \tag{31}$$

$\delta(t)$ may be as high as 1 (which is when every firm has a unique technology) and as low as $\frac{1}{m^t}$ (which is when all firms use an identical technology). Figure 9 captures the time series of $\delta(t)$ averaged over the $1,000$ replications, using the baseline parameter configuration. The "inverted S" shape of $\delta(t)$ is a general property.

The degree of technological diversity is clearly determined by the ruggedness of the landscape the firms are climbing in their search for a local optimum. Given the relationships between the values of N and K and the ruggedness of the landscape, one expects the long-run technological diversity to be a function of N and K as well. For instance, recall that $K = 0$ implies a smooth landscape with a unique optimum. When the technological environment is as simple as this, there should be a perfect convergence in the long run, since everyone is in search of the common optimum. However, for $K > 0$ – as in the baseline simulation captured in Figures 7-9 – one cannot rule out the possibility that there are multiple local optima and, hence, the firms' technological paths may diverge from one another over time as they evolve toward different optima. A thorough investigation of how N and K affect technological diversity is carried out in Section 5.

5 Comparative Dynamics

We now investigate how various parameters of the model influence the industry dynamics. The approach taken is to vary the value of one specific parameter, while holding fixed the values of other parameters at their baseline levels. The parameters of interest and their values considered for our analyses are listed in Table 2.

For each parameter configuration, a single run of the simulation will generate the time series of the endogenous variables which describe the evolving structure of the industry. In the ensuing analyses, we focus on the entry/exit measures – i.e., $\left|E^t\right|$ and $\left|L^t\right|$ – as well as the measure of intra-industry technological diversity, $\delta(t)$.

We begin by first observing the overall magnitude of the turnovers in the industry. The simplest measures to look at for this purpose are the total number of entries and the total number of exits which occurred over the entire horizon (4,000 periods) – i.e., $\sum_{t=1}^{4000}\left|E^t\right|$ and $\sum_{t=1}^{4000}\left|L^t\right|$. A high turnover rate is implied when an industry has a high total number of entries *simultaneously* with a high total number of exits. Since in our model the industry is "born" in $t = 0$ with zero incumbent, the difference between the two aggregate measures then gives the number of surviving firms at the end of the horizon. As we compute these values for each of the 1,000 replications, we obtain 1,000 independent realizations for each measure for a specific configuration of parameter values.

Tables 3(a) and 3(b) – 3(a) for (a, f, r, b, α) and 3(b) for (N, K) – report the mean values of $\sum_{t=1}^{4000}\left|E^t\right|$, $\sum_{t=1}^{4000}\left|L^t\right|$, and $\sum_{t=1}^{4000}\left|E^t\right| - \sum_{t=1}^{4000}\left|L^t\right|$ with the standard deviations reported inside the parentheses. With the sole exception of f, the total number of entries and the total number of exits move in tandem: *When an industry has a high (low) rate of entry, it also has* *a high (low) rate of exit.* Furthermore, we observe the following property with respect to the relevant parameters.

Property 1: Both the total number of entries and the total number of exits are higher – i.e., there exists a greater turnover of firms – when: 1) the market demand (a) is larger; 2) the pool of potential entrants (r) is larger; 3) the start-up fund (b) is smaller; and 4) the industry-wide search propensity (α) is lower.

We find that the impact of f on $\sum_{t=1}^{4000} \left|E^t\right|$ is non-monotonic, although $\sum_{t=1}^{4000} \left|L^t\right|$ rises monotonically in f. Both N and K have non-monotonic effects on $\sum_{t=1}^{4000} \left|E^t\right|$ and $\sum_{t=1}^{4000} \left|L^t\right|$.

The next property captures the impact of the parameters on the long-run structure the industry converges to.

Property 2: The number of surviving incumbents in the long-run is higher when: 1) the market demand (a) is larger; 2) the fixed cost (f) is lower; 3) the pool of potential entrants (r) is larger; and 4) the search space (N) is larger.

The impacts of b, α, and K on the net entry measure, are non-monotonic and, thus, necessitate further investigation.

In the following subsections, we further refine and reinforce these results by investigating in detail the time paths of $\left|E^t\right|$, $\left|L^t\right|$, and m^t, as well as that of $\delta(t)$. In particular, the results on $\delta(t)$ are presented formally in Properties 3-5. Note that we perform 1,000 independent replications for each configuration of parameter values. This means that we have 1,000 separate realizations of these measures for each time period. We take a simple average of these realizations for each t, thereby generating a time path of the mean values. Such time paths are generated for different parameter values and then plotted for visual comparison in Figures 10 through 16.

5.1 Effects of Demand, Cost, and Entry Conditions

We begin with the effects of a, f, r, and b. The short-run impacts of these parameters can be inferred from the static equilibrium model of Cournot oligopoly. For instance, note from equations (11), (13), and (14) that the size of the market demand, a, increases the market price, firm output (and, hence, the industry outputs), and the firm profit, holding the number of firms constant. The size of the fixed cost, f, on the other hand, has no impact on the price, or the firm's output, but has a negative impact on the firm's profit, again, holding the number of firms constant. The entry/exit tendencies in the industry can then be indirectly inferred from the equilibrium levels of firm profitabilities. Of course, in the dynamic model of firm competition in which entry and exit can take place continually, it is not at all clear how these parameters would affect the evolution of the industry structure over time.

How does the size of market demand affect the flow of firms in and out of the industry? In Figure 10, we first explore the impact of a on the number of entries which take place over time. Clearly, an increase in a raises the number of entries that take place before they cease eventually. Furthermore, the time series on the number of exits (second figure in Figure 10) show that the shakeout process is delayed and weakened in its magnitude as the size of market demand is greater. This implies that the increased number of new entrants under a higher value of "a" also tend to stay in the industry, thereby leading to a long-run and permanent expansion of the industry. The time series on m^t, the number of incumbent firms (third figure), do confirm this: there is a general expansion of the industry in the long run as the size of market demand increases.

The bottom figure in Figure 10 shows what happens to the degree of technological diversity as the market demand becomes larger. The mean technological diversity (averaged over 1,000 replications) starts from 1, where all initial entrants have heterogeneous technologies, and then gradually declines over time. However, the decline is much faster in a market with a higher value of a: The mean diversity decreases in a. Why is this so? This is due to the increase in the number of firms, m^t, resulting from an increase in a. Note that the technological environment is uniquely defined by N and K which are held fixed for all values of a considered here. This implies that the number of locally optimal technologies (averaged over 1,000 replications) tends to be constant for different values of a. However, in a market with a higher value of a, a greater number of firms enter the market and the long-run searches by many of these firms lead them toward common technologies. Whether the markets are more or less concentrated, the number of groups of firms sharing common technologies – technology clusters mainly determined by N and K – tends to be fixed. Since the size of the membership in each of these groups then increases in a, the mean technological diversity will decline in the population size.

Figure 11 captures the effects of fixed cost f on firm turnovers. At the initial glance, the impact of f on the number of entries over time is somewhat counter-intuitive: one would have expected the number of entries to be negatively related to fixed cost, since it is more difficult to make profits in a market with a higher value of f – this is what the static equilibrium model of Cournot oligopoly would predict. Surprisingly, the time series in Figure 11 show that the size of f has little impact on the number of entries. This result is driven by the assumption in this model that the decision to enter the market is based on the consideration of relative production efficiencies rather than that of expected profits. On the other hand, the size of f must have an impact on the number of exits, since the decision to exit the market in this model is based on the realized profits (through the firm's wealth balance). The middle figure on the time series of the number of exits clearly shows that this is the case. In a market with a higher value of f, the shakeout process is accelerated and its impact is much more significant. Combining the entry dynamic with the exit dynamic, we then find that the number of operating firms is reduced in the long run for higher values of f: A market with a higher fixed cost tends to be more concentrated.

Figure 11 also shows that the mean technological diversity is positively related to the size of fixed cost. This result is driven by the same intuition as that which drove the influence of a on diversity. As f rises in value, the number of operating firms declines. This reduces the number of firms sharing identical technologies and increases the overall diversity.

The impact the size of potential entrants pool (r) has on the industry dynamics is explored in Figure 12. How large the pool of potential entrants is depends on the attractiveness of the given industry relative to other industries as well as the transferability of production knowledge from one industry to the next. If the industries have similar production systems where the relevant sets of knowledge overlap to a significant extent, then it seems reasonable that the size of the potential entrants pool can be quite large. How does the magnitude of r then affect the industry dynamics? Do industries facing different sizes of entrant pools behave differently over time? Figure 12 shows that r increases the number of entries as well as the number of exits significantly in the short run, which implies that the turnover is likely to be higher in those industries with higher values of r. For when the industries differ in terms of the availability of potential entrants, our result is then fully consistent with the empirical findings that there exists a high correlation between the entry and exit rates across industries. Finally, the number of incumbent firms is greater for all t when r is higher.

The degree of technological diversity decreases in r. An industry with a bigger pool of potential entrants tends to have a lower degree of diversity as the increasing number of incumbent firms converge on a limited number of locally optimal technologies.

It is shown in Figure 13 that the size of the start-up fund has little impact on the number of entries, though it appears to have a (weak) positive effect in the short run and a (weak) negative effect in the long run. However, in terms of the exit dynamics, a larger amount of start-up fund systematically delays the shakeout process by allowing the inefficient firms to linger on for a longer period of time. This, of course, increases the number of incumbents in the short run, but it has a negligible impact in the long run. Once all the exits have taken place, the resulting market concentration is unlikely to be affected by the size of b.

The degree of technological diversity is relatively insensitive to b, though one detects a weakly negative relationship between them. A market with a greater b tends to accommodate a larger number of firms in the short run and this leads to a decrease in the mean diversity.

The results on the mean technological diversity, involving the demand/cost/entry parameters, may then be summarized as follows.

Property 3: The degree of technological diversity is higher in the long run when: 1) the market demand is smaller; 2) the fixed cost is higher; 3) the potential entrants pool is smaller; and 4) the start-up fund is smaller.

5.2 Effects of Search Propensity (α)

The intensity of a firm's innovation/imitation activity is captured by α, the probability of search in each period. If different industries have different degrees of search propensity, will their evolutionary dynamics be influenced by such differences?

Figure 14 shows that the industry-wide propensity to search does indeed have a significant impact on the dynamics of the market structure. There is a general decline in the number of entries (in the short- to medium-run) as the overall search activity in an industry is more intense – this is intuitive as the incumbent firms' active search (and the consequent increase in the average efficiency level) will discourage potential entrants from entering the market. There is also a decline in the number of exits during the shakeout phase – more intensive search activities by all incumbent firms make it less likely that any given firm will remain a laggard. As these two effects combine together, a higher value of α appears to lower the number of incumbents in the market. This then implies that an industry with firms having a greater propensity to search tends to experience less turnover of firms – the select few who enter the market have a greater likelihood of remaining in the market.

The degree of technological diversity captured in Figure 14 tells an interesting story. As search becomes more intensive, the degree of diversity declines over time. Notice that the mean technological diversity remains quite high throughout the entire horizon for $\alpha = 0$. When $\alpha = 0$, there is no innovation or imitation. Any change in the distribution of the chosen technologies is purely due to the market selection imposed on the diverse efficiency levels that the firms are endowed with at the time of their entry. As the general level of efficiency in the industry reaches a sufficiently high level in the long run, further entry becomes rare and we tend to see a highly stable market with the surviving firms still operating with the diverse technologies that they originally entered the industry with. There is, hence, no technological convergence, when there are no innovations or imitations. As α rises above zero, however, more active search by the incumbent firms drive them toward a group of common technologies, thereby reducing the overall technological diversity in the long run.

Property 4: The degree of technological diversity is higher when the firms have weaker propensity to search for improved technologies.

5.3 Effects of Technological Complexity

Industries may be differentiated in terms of how complex their production processes are. In modern manufacturing where many of the goods and services are assembled, the degree of technological complexity is often determined by the nature of the product that is generated (or put together) at the end of the process:

> [T]he production process for an automobile engine (just the engine!) had 130 steps: drill, turn, drill, grind, attach component, transfer to next station, and so forth ... The most archaic plate glass-making process, by contrast, had only basic five production steps. — Utterback (1994), p.118

Some industries are then characterized by production processes that entail a large number of component activities, while others involve processes with only a few activities. Furthermore, a subset of these component activities may be inter-dependent with one another to such an extent that there are technological complementarities that make the search for the optimal process design non-trivial.

In our model, these aspects of the production process are captured using N and K, where N is the total number of component activities in the production system and K is the number of other activities that are coupled to each activity. Thus far, we have kept the values of N and K constant at 16 and 2, respectively. A relevant question to ask in this context is how two industries having different degrees of technological complexity – i.e., different values of N or K – will evolve over time. Note that a higher value of N for an industry implies that the firms must look for an improved technology in a larger (higher dimensional) search space, while a higher value of K implies that a larger subset of activities are coupled to each other and, thus, the technology landscape is more rugged.

As shown in Kauffman (1993), an increase in either or both of these parameters raises the number of local optima on the technology landscape. In this section, I explore how the size of the search space and the ruggedness of the technology landscape affect the speed and the effectiveness of the firm-level search and what this implies in terms of the evolving industry structure.

5.3.1 Size of the Search Space (N)

Let us start with the impact of N on the mean technological diversity [see the bottom figure in Figure 15]. In the short run when the firms are actively searching for the (locally) optimal technology, the expansion of the search space – i.e., an increase in N – unambiguously raises the initial extent of diversity as there exists a wider variety of component activities over which the new entrants may start out with unequal production methods. In addition, the average gain in efficiency realized from a successful idea tends to be smaller when N is larger and, hence, the overall process of search tends to slow down as well. This further adds to the initial gain in diversity coming from a higher value of N. A more fundamental result, however, is that the higher degree of technological diversity (introduced into the industry through a higher value of N) tends to be permanent. As shown in Figure 15, the diversity measure stays high even in the long run when N is large. The intuition lies in the fact that a higher value of N raises the number of local optima on the technology landscape so that the incumbent firms are all likely to move toward *distinct* optima (technologies) in the long run: The probability of any two firms sharing a common technology in the long run steady-state is lower when N is larger.

How does N affect the size of entry and exit? We find that the effects that N has on both $\left|E^t\right|$ and $\left|L^t\right|$ tend to change their signs over time. Both $\left|E^t\right|$ and $\left|L^t\right|$ increase in N in the short run (between $t = 1$ and $t \approx 40$). This is clearly captured in Figure 15. For $t > 40$, however, the effects reverse their signs such that $\left|E^t\right|$ and $\left|L^t\right|$ both *decrease* in N. Because this long-run impact of N is not clearly visible in Figure 15, I provide additional evidence in Table 4 where the steady-state numbers of entries and exits are reported as the per-period average over the last 2,000 periods from $t = 2001$ to $t = 4000$. It is clear that the mean numbers of entries and exits monotonically decline in N in the long run.

Why do we observe these effects and what causes their reversal over time? We first focus on the short-run result that both $\left|E^t\right|$ and $\left|L^t\right|$ monotonically increase in N. Given the fixed rate of idea generation, α, and the fact that each idea entails changing the method in only *one* of the component activities, a search space with a greater number of dimensions (N) implies that the average gain in efficiency from a successful idea is relatively smaller. At the population level, the distribution of the efficiency levels held by the active firms in the market tends to shift up at a much slower rate when N is large. Since the entry decision of a potential entrant is based on the efficiency level of the minimally efficient active firm, there then exists a larger room for the potential entrants to come into the market at any given point in time. The range of the entrants' efficiency levels will also be wider and this implies that the ensuing competition will drive a larger number of incumbent firms out of the market, thereby amplifying the shakeout process in the beginning of the horizon. The numbers of entries and exits captured in Figure 15 for $0 < t < 40$ confirm these intuitions: N raises the number of entries in the short run, but it also raises the number of exits that follow. This result then implies that an industry with a production system which consists of a greater number of component activities is likely to experience a greater rate of firm turnovers in the short run. Furthermore, a higher N appears to initially promote entry to a greater extent than it does exit, so that the number of incumbents in the market increases in N in the short run.

In contrast to the short-run relationship, the long-run relationship between N and $\left|E^t\right|$ (or $\left|L^t\right|$) is negative. The intuition behind this result is based on the fact that the long-run time frame considered here corresponds to that in which many incumbent firms have already attained their local optima. Note that the potential entrants do not come into the market unless they have a technology which is at least as efficient as the one held by the least efficient incumbent. Clearly, the probability of a potential entrant getting a sufficiently good draw on the technology is lower if a larger number of incumbent firms are already at the local optima – the entrant must pick a technology which is superior to the ones that are found to be locally optimal by the incumbent firms. Since a larger N raises the number of local optima, this tends to also raise the proportion of incumbents who are already at various local optima in any given period during the long-run time frame. A larger N then reduces the number of entries and the number of exits that follow. In the end, the long-run number of surviving incumbents tends to be lower for a larger value of N.

5.3.2 Inter-activity Complementarity (K)

For a given N, an increase in K raises the number of local optima without an expansion of the search space. For a randomly chosen firm with its current technology, the expected number of trials to get to a local optimum is then lower when K is higher. Consequently, potential entrants find it more difficult to enter the market, since the incumbents are likely to find their local optima rather quickly. Unless the potential entrants enter the market with a technology which represents a superior local optimum, it is not likely that they will enter the market. Since only the relatively efficient firms will be able to enter in this case, the number of exits will be low as well. As confirmed in Figure 16, the rate of firm turnovers then declines in K.

The mean technological diversity rises in K [bottom figure in Figure 16]. This property is as expected: an increase in K and the consequent increase in the number of local optima raises the probability that each individual firm will have a unique technology in the long run.

Combining together the effects of N and K on the mean technological diversity, we establish the following property.

Property 5: The degree of technological diversity is higher when 1) the production process entails a greater number of component activities and 2) there is a greater degree of interdependence among component activities.

6 Concluding Remarks

The model presented here is based on the view that the production process is a system of inter-dependent activities and the firm is an adaptive entity whose survival depends on its ability to discover ways to perform various activities with greater efficiency than its rivals. The search for improved technology is conducted on a pre-defined technological landscape, so that the specifications of a firm's current knowledge and its path of knowledge acquisition are made explicit. The evolution of the industry, as well as the technology, is then driven by the innovation/imitation process the firms use and the attributes of the market in which the selection pressure is continually applied on the population of firms through the entry of new firms and the consequent competition among the incumbent firms.

Many well-known empirical regularities arose naturally in this model, including those on firm turnovers, evolving market structure, and the intra-industry technological diversity. In addition to the generation of regularities at the industry level, we performed extensive comparative dynamics analyses in order to investigate how these regularities are affected by various industry-specific factors such as the attributes of the market environment, search propensities, and the nature of the technology space in which individual firm's learning takes place. All of the industry-specific factors considered here are found to have significant impacts on the shakeout dynamics of the industry as well as the evolution of the intra-industry technological diversity. Of particular interest are the results concerning the turnover of the firms. We find that the turnover is higher – i.e., the aggregate numbers of entries and exits over time are simultaneously greater – when: 1) the market demand is larger; 2) the potential entrants pool is larger; 3) the start-up fund is smaller; and 4) the firms in the industry have a lower propensity to search. The number of surviving incumbents is higher in the long run when: 1) market demand is larger; 2) the fixed cost is lower; 3) the pool of potential entrants is larger; and 4) the production process entails a greater number of component activities. Finally, the persistent technological diversity within a given industry is found to depend on the extent of complexity in the production process which gives rise to multiplicity of optimal technologies. Furthermore, the degree of technological diversity is observed to be higher in the long run when: 1) the market demand is smaller; 2) the fixed cost is higher; 3) the potential entrants pool is smaller; 4) the start-up fund is smaller; 5) the firms have weaker propensity to search for improved technologies; 6) the production process entails a greater number of component activities; and 7) there is a greater degree of interdependence among component activities.

There are two ways in which this model can be enriched in a substantive manner. First, innovations and imitations are modelled here as being purely serendipitous – the firms do not make conscious investment decisions to engage in research and development to raise the probability of generating a useful idea. Incorporating such an investment decision into the current model will be an important step. An issue that is critical to pursuing this extension is whether to model the investment decision as being driven by foresight on the part of the firms or by a reinforcement learning mechanism which relies purely on their past experiences.

The second aspect of the current model which could be enriched involves the technological environment within which the firm search takes place. In the current model, the firms search in a static technological environment in that the mapping between the methods vector and the efficiency level remains fixed over time. But what if the environment is subject to continual external shocks? How would the shakeout phenomenon and the evolving industry dynamic be affected by the extent of such volatility in production efficiency? This issue can be tackled by allowing the efficiency contributions of the individual methods to be determined fresh each period in a systematic manner. By controlling the way the new efficiency level for a method deviates from the current efficiency level, one can capture the degree of volatility in the technological environment.

A final remark that I would like to make is that the current model can also be used to explore the empirically relevant issues other than the ones considered in this paper. Some examples include the evolving distribution of the firm sizes over the course of industry development as well as the age distribution of the surviving firms and the relationship between a firm's age and the rate of survival. An agent-based computational model such as the one presented here provides a natural laboratory in which such questions can be answered through a series of controlled computational experiments. Work is currently underway to address these issues.

References

[1] Bergstrom, T. C. and H. R. Varian, "When Are Nash Equilibria Independent of the Distribution of Agents' Characteristics?," *Review of Economic Studies*, 52 (1985), 715-718.

[2] Carroll, G. R. and M. T. Hannan, *The Demography of Corporations and Industries*, Princeton University Press, Princeton, New Jersey (2000).

[3] Caves, R. E., "Industrial Organization and New Findings on the Turnover and Mobility of Firms," *Journal of Economic Literature*, XXXVI (1998), 1947-1982.

[4] Dawid, H., "Agent-Based Models of Innovation and Technological Change," in: Tesfatsion, L. and K. L. Judd (editors), *Handbook of Computational Economics, Volume 2: Agent-Based Computational Economics*, Handbooks in Economics Series, North Holland, 2006.

[5] Dunne, T., M. J. Roberts, and L. Samuelson, "Dynamic Patterns of Firm Entry, Exit, and Growth," *RAND Journal of Economics*, 19 (1988), 495-515.

[6] Epstein, J. M., "Remarks on the Foundations of Agent-Based Generative Social Science," in: Tesfatsion, L. and K. L. Judd (editors), *Handbook of Computational Economics, Volume 2: Agent-Based Computational Economics*, Handbooks in Economics Series, North Holland, 2006.

[7] Ethiraj, S. K. and D. Levinthal, "Modularity and Innovation in Complex Systems," *Management Science*, 50 (2004), 159-173.

[8] Ghemawat, P., "Competitive Advantage and Internal Organization: Nucor Revisited," *Journal of Economics & Management Strategy*, 3 (1995), 685-717.

[9] Gort, M. and S. Klepper, "Time Paths in the Diffusion of Product Innovations," *Economic Journal*, 92 (1982), 630-653.

[10] Ichniowski, C., K. Shaw, and G. Prennushi, "The Effects of Human Resource Management Practices on Productivity: A Study of Steel Finishing Lines," *American Economic Review*, 87 (1997), 291-313.

[11] Jovanovic, B. and G. M. MacDonald, "The Life Cycle of a Competitive Industry," *Journal of Political Economy*, 102 (1994), 322-347.

[12] Kauffman, S. A., 1993, *The Origins of Order*, Oxford University Press, Oxford.

[13] Klepper, S., "Firm Survival and the Evolution of Oligopoly," *RAND Journal of Economics*, 33 (2002), 37-61.

[14] Klepper, S. and E. Graddy, "The Evolution of New Industries and the Determinants of Market Structure," *RAND Journal of Economics*, 21 (1990), 27-44.

[15] Klepper, S. and K. L. Simons, "Technological Extinctions of Industrial Firms: An Inquiry into their Nature and Causes," *Industrial and Corporate Change*, 6 (1997), 379-460.

[16] Klepper, S. and K. L. Simons, "Dominance by Birthright: Entry of Prior Radio Producers and Competitive Ramifications in the U.S. Television Receiver Industry," *Strategic Management Journal*, 21 (2000a), 997-1016.

[17] Klepper, S. and K. L. Simons, "The Making of an Oligopoly: Firm Survival and Technological Change in the Evolution of the U.S. Tire Industry," *Journal of Political Economy*, 108 (2000b), 728-760.

[18] Levinthal, D., "Adaptation on Rugged Landscapes," *Management Science*, 43 (1997), 934-950.

[19] MacDuffie, J. P., "Human Resource Bundles and Manufacturing Performance: Organizational Logic and Flexible Production Systems in the World Auto Industry," *Industrial and Labor Relations Review*, 48 (1995), 197-221.

[20] MacDuffie, J. P. and J. F. Krafcik, "Integrating Technology and Human Resources for High-Performance Manufacturing: Evidence from the International Auto Industry," in: T. A. Kochan and M. Useem, eds., *Transforming Organizations* (Oxford University Press, Oxford), 1992.

[21] Milgrom, P. and J. Roberts, "The Economics of Modern Manufacturing: Technology, Strategy, and Organization," *American Economic Review*, 80 (1990), 511-528.

[22] Milgrom, P. and J. Roberts, "Complementarities and Fit: Strategy, Structure, and Organizational Change in Manufacturing," *Journal of Accounting and Economics*, 19 (1995), 179-208.

[23] Milgrom, P., Y. Qian, and J. Roberts, "Complementarities, Momentum, and the Evolution of Modern Manufacturing," *American Economic Review*, 81 (1991), 84-88.

[24] Porter, M., "What Is Strategy?," *Harvard Business Review*, November-December 1996, 61-78.

[25] Rivkin, J., "Imitation of Complex Strategies," *Management Science*, 46 (2000), 824-844.

[26] Rivkin, J. and N. Siggelkow, "Balancing Search and Stability: Interdependencies among Elements of Organizational Design," *Management Science*, 49 (2003), 290-311.

[27] Sutton, J., "Gibrat's Legacy," *Journal of Economic Literature*, XXXV (1997), 40-59.

[28] Tesfatsion, L. and K. L. Judd (editors), *Handbook of Computational Economics, Volume 2: Agent-Based Computational Economics*, Handbooks in Economics Series, North Holland, 2006.

[29] Utterback, J. M., *Mastering the Dynamics of Innovation: How Companies Can Seize Opportunities in the Face of Technological Change*, Harvard Business School Press, Boston, Massachusetts, 1994.

[30] Vives, X., "Complementarities and Games: New Developments," *Journal of Economic Literature*, XLIII (2005), 437-479.

Table 1: Set Notations

Notation	Definition
S^t	Set of surviving firms at the end of t
S_a^t	Those in S^t which were active in t
S_{-a}^t	Those in S^t which were inactive in t
R^t	Set of potential entrants at the beginning of t
E^t	Set of actual entrants in t
M^t	Set of firms poised to compete in t $(= S^{t-1} \cup E^t)$
S_b^t	Those in S^t which were profitable in t
L^t	Set of firms which exit the industry at the end of t

Table 2: List of Parameters and Their Values

Notation	Definition	Baseline Value	Parameter Values Considered
N	No. of Activities	16	{10, 16, 22, 32}
K	Degree of Complexity	2	{1, 2, 4, 6}
r	No. of Potential Entrants Per Period	10	{5, 10, 20, 40}
f	Fixed Cost	20	{5, 10, 20, 40}
a	Market Size (Demand Intercept)	200	{100, 200, 400, 600}
b	Start-Up Budget for a New Entrant	100	{0, 10, 50, 100, 200}
d	Threshold Wealth Balance for Exit	0	0
α	Probability of Search	1.0	{0, .2, .4, .6, .8, 1}
T	Time Horizon	4,000	4,000

Table 3 (a): Total Entrants, Total Exits, and the Net Entrants*

Parameter	Values	$\sum_{t=1}^{4000} \lvert E^t \rvert$	$\sum_{t=1}^{4000} \lvert L^t \rvert$	$\sum_{t=1}^{4000} \lvert E^t \rvert - \sum_{t=1}^{4000} \lvert L^t \rvert$
a	100	52.938 (14.5676)	41.405 (13.0212)	11.533 (3.02491)
	200	98.55 (24.8333)	71.842 (21.4587)	26.708 (6.35351)
	400	163.011 (42.7108)	108.507 (34.7703)	54.504 (13.0733)
	600	219.33 (61.5639)	137.537 (49.705)	81.793 (19.2965)
f	5	97.598 (26.8411)	55.729 (19.4411)	41.869 (12.5571)
	10	98.24 (26.4372)	64.878 (20.78)	33.362 (9.10549)
	20	98.55 (24.8333)	71.842 (21.4587)	26.708 (6.35351)
	40	94.636 (20.8792)	74.211 (19.0068)	20.425 (4.11963)
r	5	70.099 (18.7419)	45.327 (15.4612)	24.772 (5.9531)
	10	98.55 (24.8333)	71.842 (21.4587)	26.708 (6.35351)
	20	133.036 (32.2883)	105.017 (29.2701)	28.019 (6.05072)
	40	187.051 (49.934)	157.157 (47.356)	29.894 (6.0124)
b	0	107.438 (30.2537)	81.541 (26.6077)	25.897 (6.36938)
	10	104.867 (28.4903)	78.788 (25.049)	26.079 (6.15846)
	50	100.533 (25.8006)	74.333 (22.0036)	26.2 (6.27514)
	100	98.55 (24.8333)	71.842 (21.4587)	26.708 (6.35351)
	200	95.168 (24.9625)	68.527 (21.7563)	26.641 (6.28853)
α	0.0	336.807 (51.6723)	304.134 (50.1996)	32.673 (3.43014)
	0.2	165.139 (36.9995)	138.266 (34.4218)	26.873 (5.59537)
	0.4	134.764 (33.6389)	108.152 (30.6356)	26.612 (5.77848)
	0.6	115.361 (27.3696)	88.745 (24.262)	26.616 (6.10014)
	0.8	106.001 (26.4235)	79.439 (23.0269)	26.562 (6.20473)
	1.0	98.55 (24.8333)	71.842 (21.4587)	26.708 (6.35351)

* Mean over 1,000 replications. Standard Deviations are provided in the parenthesis.

Table 3 (b): Total Entrants, Total Exits, and the Net Entrants*

Parameter	Values	$\sum_{t=1}^{4000} \lvert E^t \rvert$	$\sum_{t=1}^{4000} \lvert L^t \rvert$	$\sum_{t=1}^{4000} \lvert E^t \rvert - \sum_{t=1}^{4000} \lvert L^t \rvert$
N	10	153.492 (40.8841)	114.204 (41.7175)	39.288 (2.98294)
	16	98.55 (24.8333)	71.842 (21.4587)	26.708 (6.35351)
	22	93.359 (24.405)	70.61 (22.2903)	22.749 (5.01509)
	32	95.304 (22.6531)	75.293 (21.2315)	20.011 (3.81043)
K	1	113.405 (61.9597)	86.9 (59.6908)	26.505 (5.88325)
	2	98.55 (24.8333)	71.842 (21.4587)	26.708 (6.35351)
	4	103.194 (18.1872)	77.514 (15.7142)	25.68 (6.16925)
	6	117.168 (18.2084)	91.769 (15.7348)	25.399 (5.94653)

* Mean over 1,000 replications. Standard Deviations are provided in the parenthesis.

Table 4: Steady-State Number of Entrants and Exits

		Mean Value per Period	
		$\frac{1}{2000}\sum_{t=2001}^{4000}\left\|E^t\right\| =$	$\frac{1}{2000}\sum_{t=2001}^{4000}\left\|L^t\right\| =$
N	10	0.012942	0.011539
	16	0.002344	0.001750
	22	0.000833	0.000520
	32	0.000131	0.000088

Figure 1: Decision Stages in Period *t*

Stage 1
Entry Decisions
(by all potential entrants in *t*)

↓

Stage 2
Innovation / Imitation Decisions
(by all surviving incumbents from *t* – *1*)

↓

Stage 3
Output Decisions and Market Competition
(by all operating firms in *t*)

↓

Stage 4
Exit Decisions
(by all operating firms in *t*)

Figure 2: Jovanovic-MacDonald Data

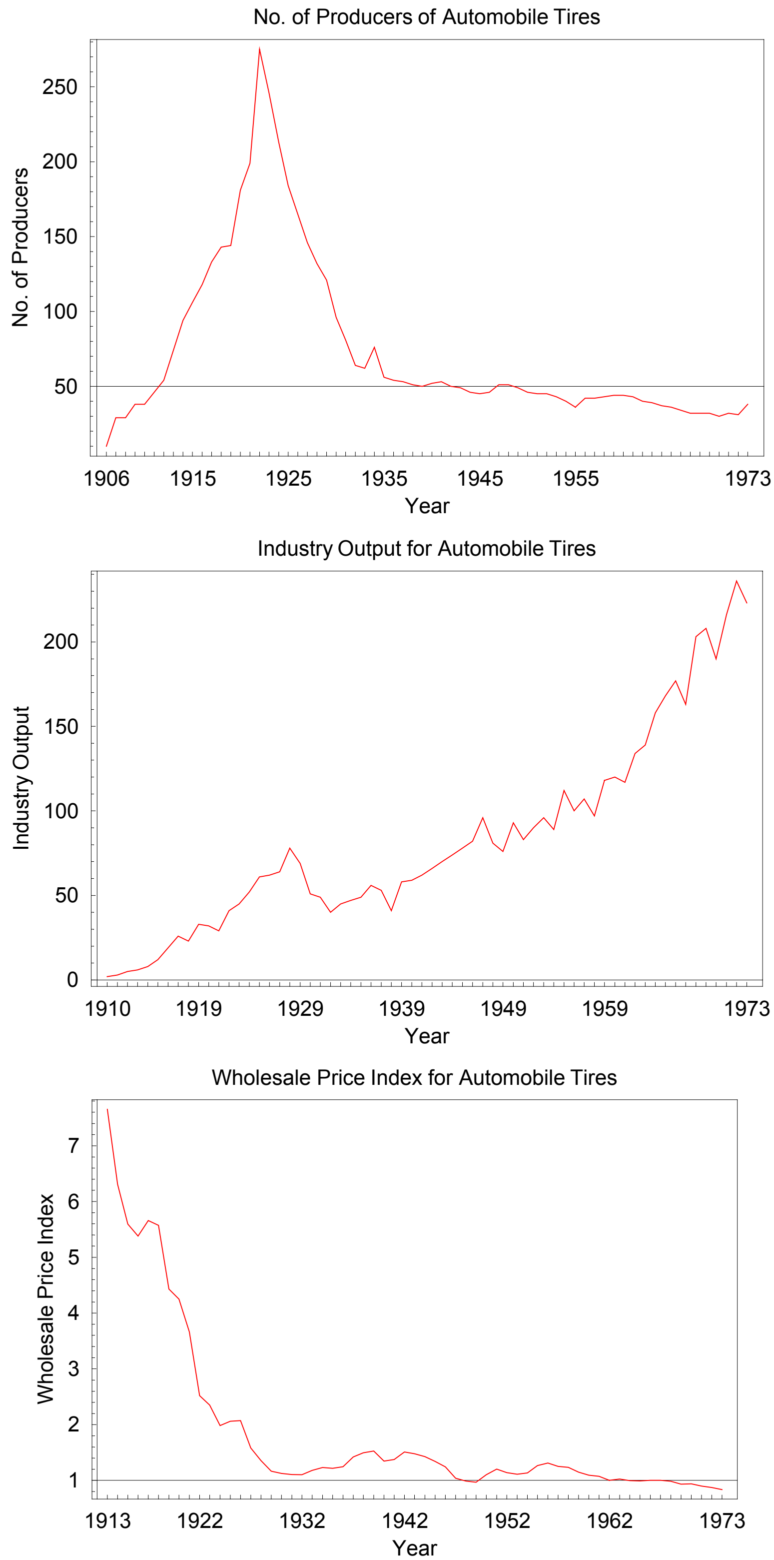

Figure 3: Baseline Results over J-M Time Span

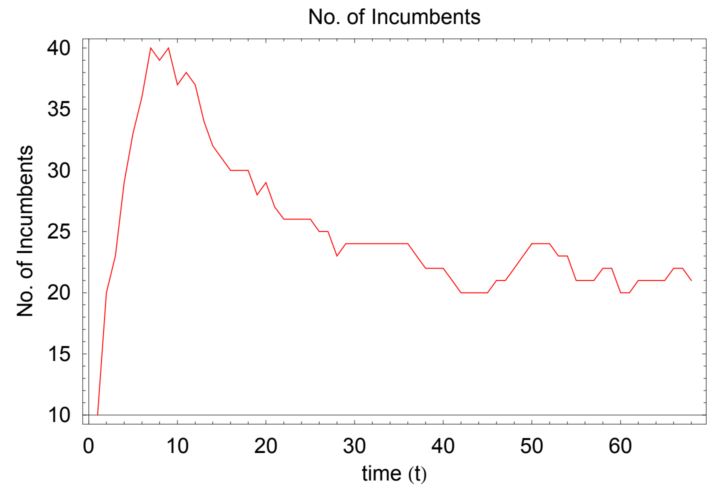

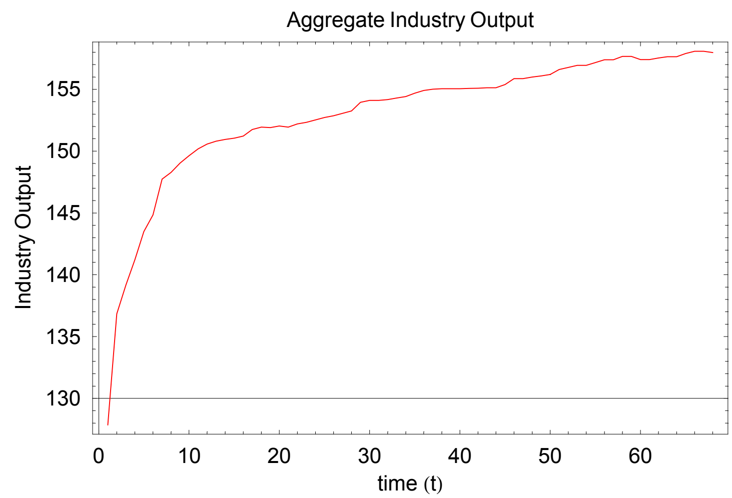

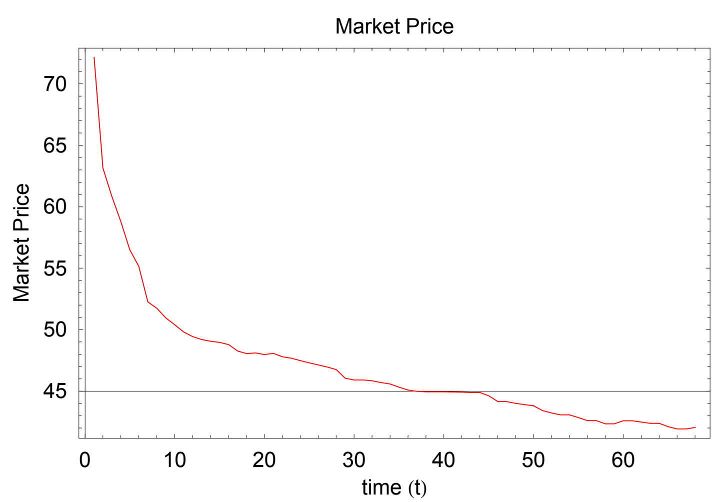

Figure 4: Turnover of Firms

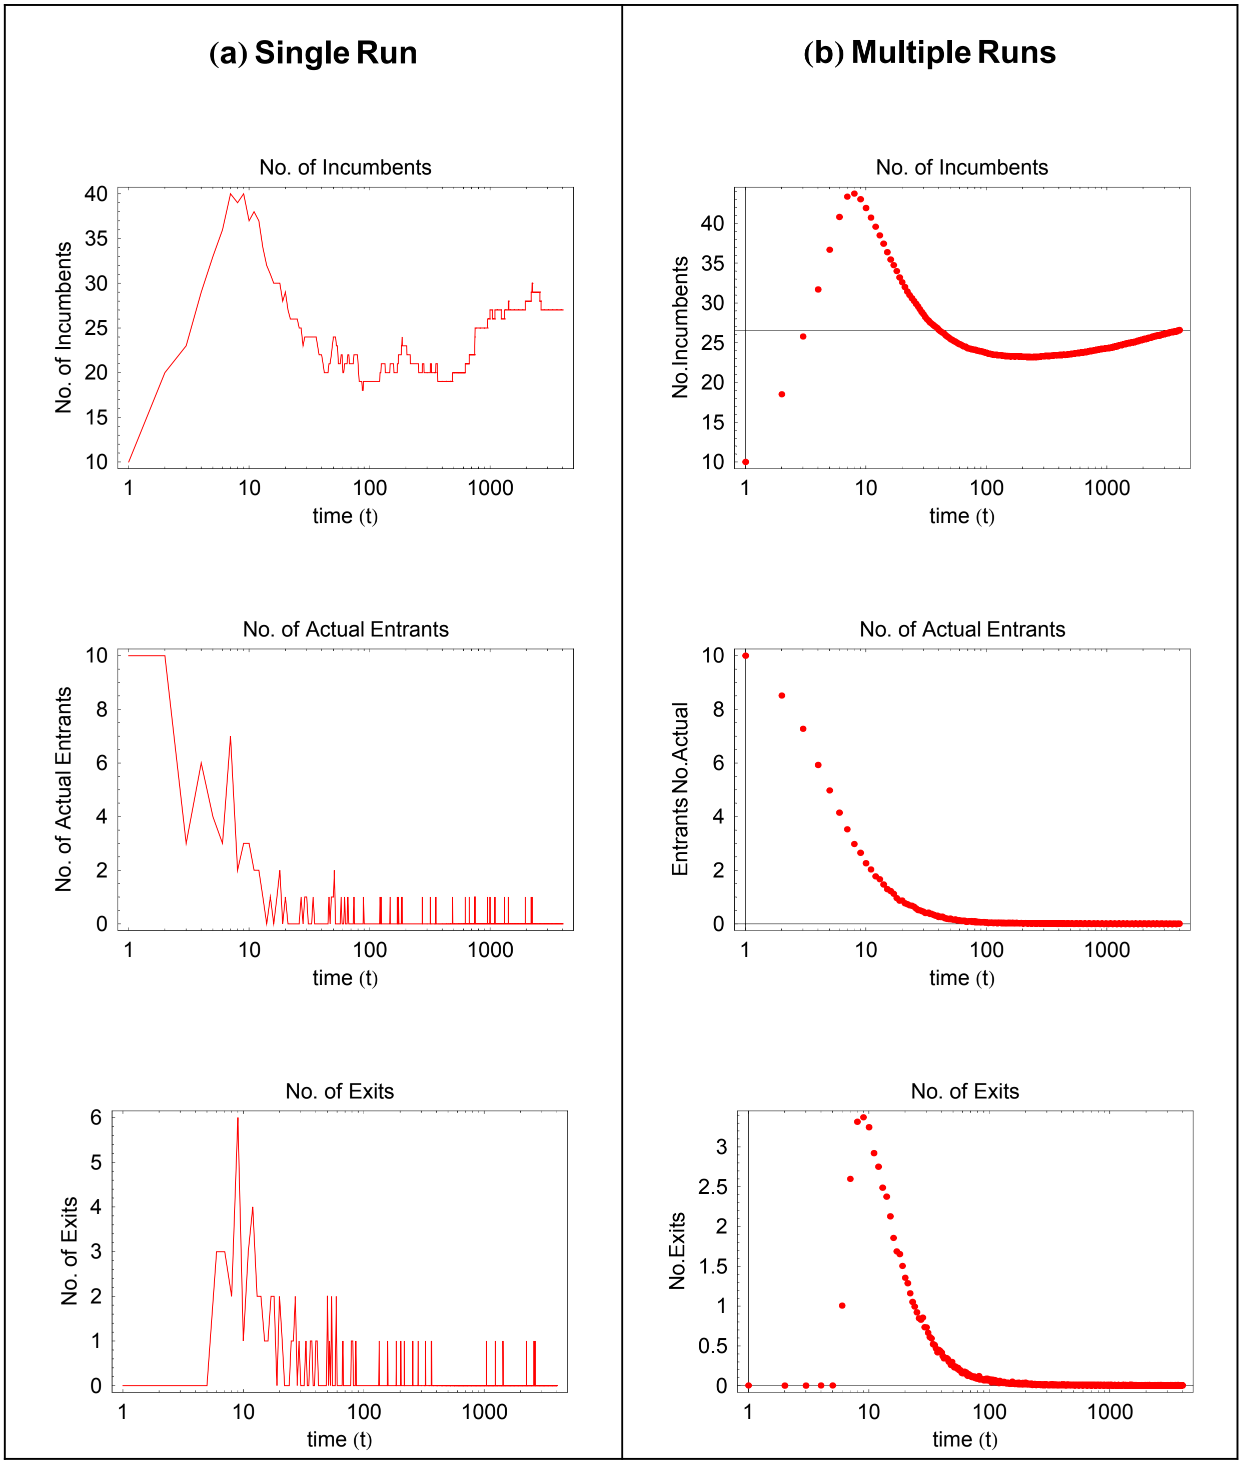

Figure 5: Market Price, Industry Output, and Concentration

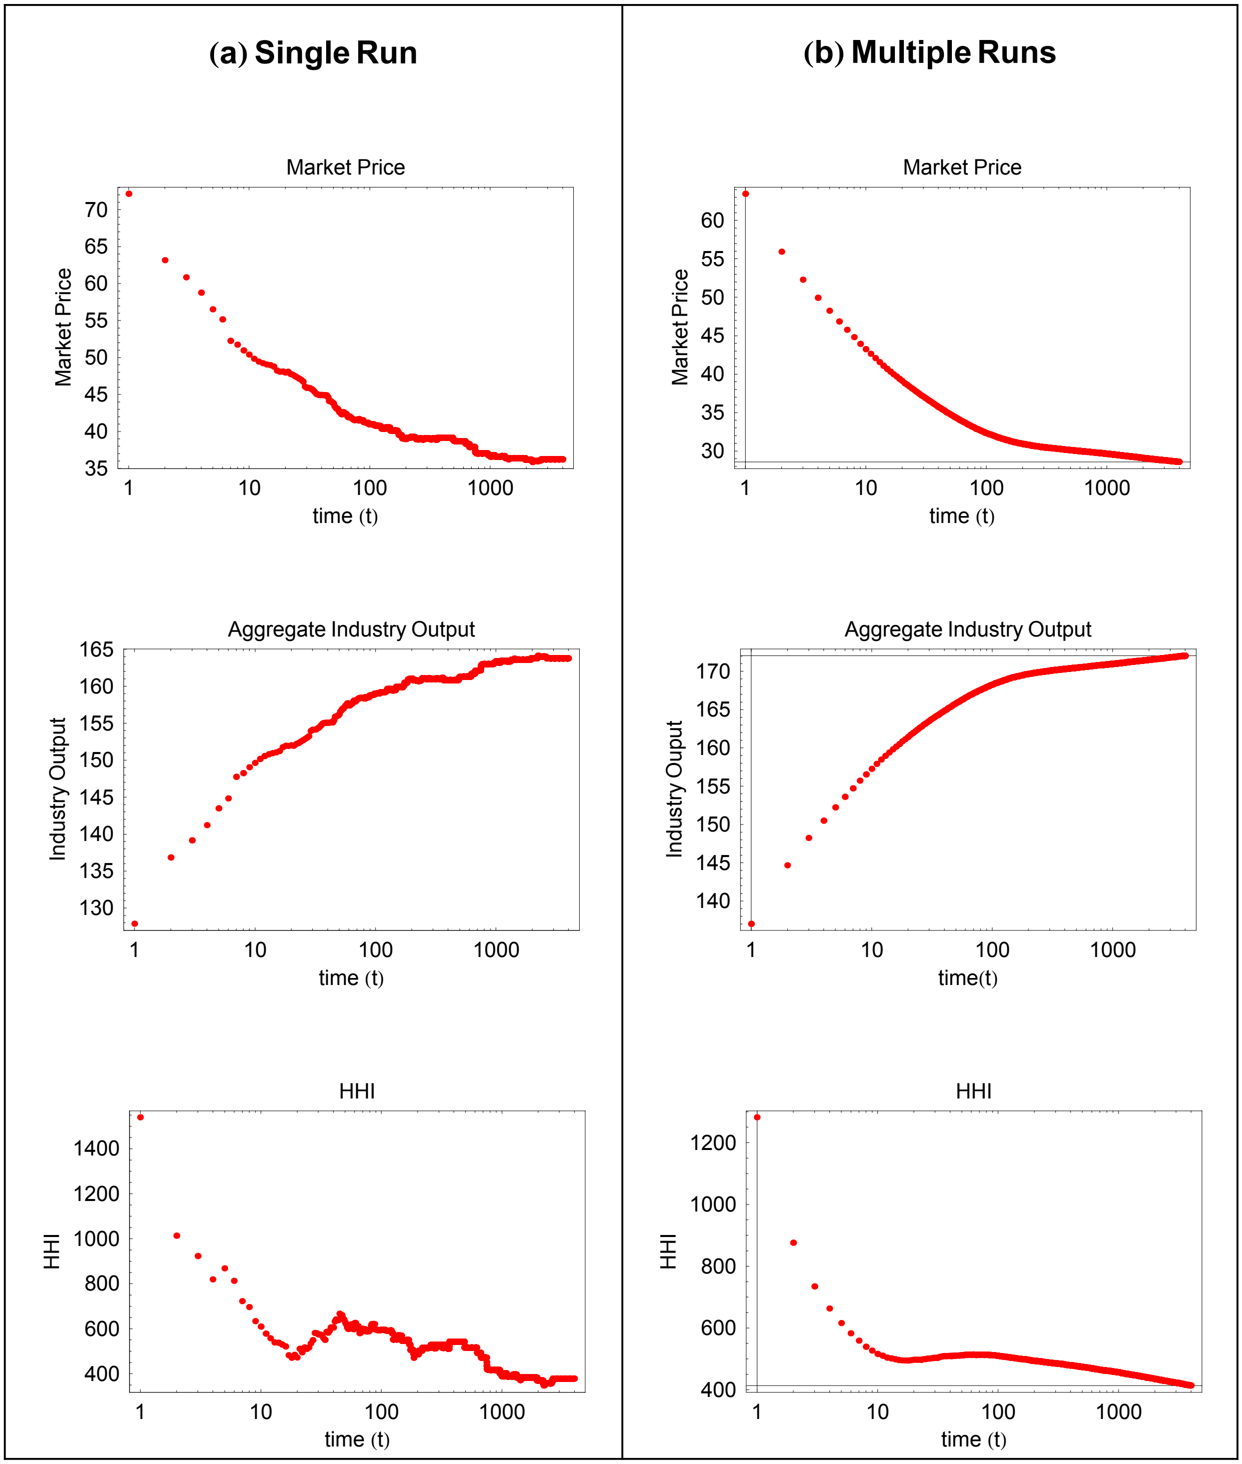

Figure 6: Distributions of Marginal Costs and Firm Outputs

(a)

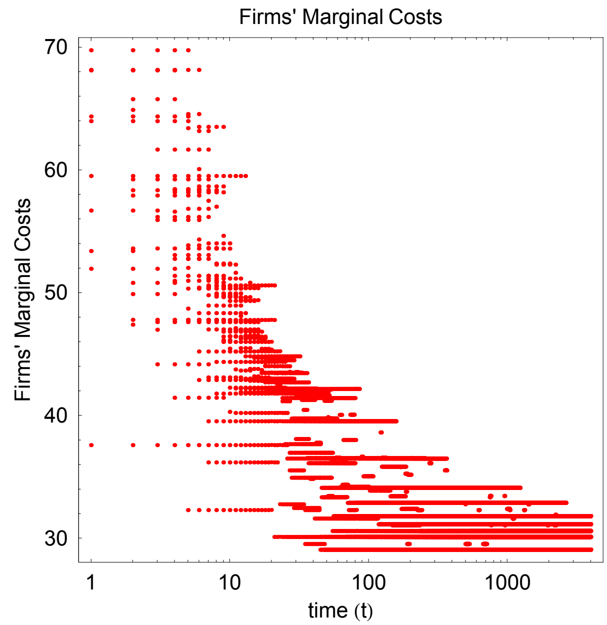

(b)

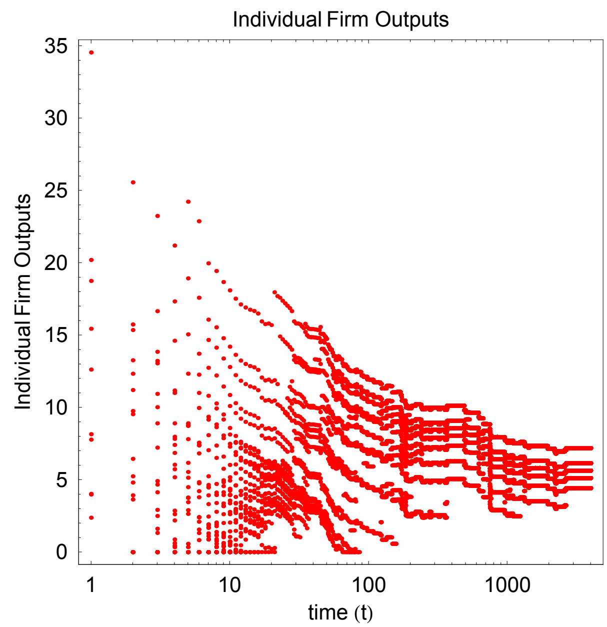

Figure 7: Number of Distinct Technologies

(a) Single Run

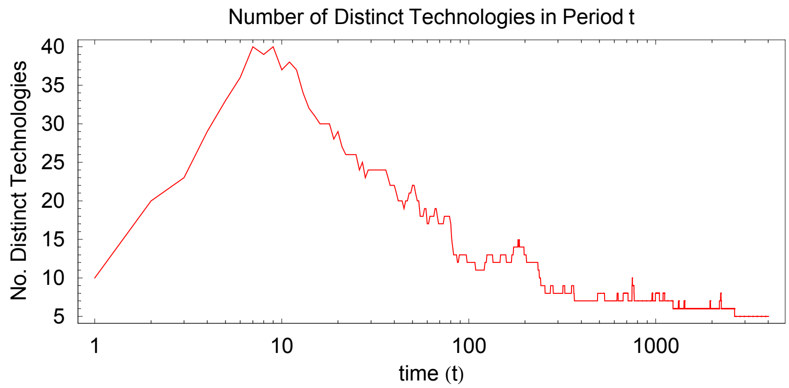

(b) 1, 000 Runs

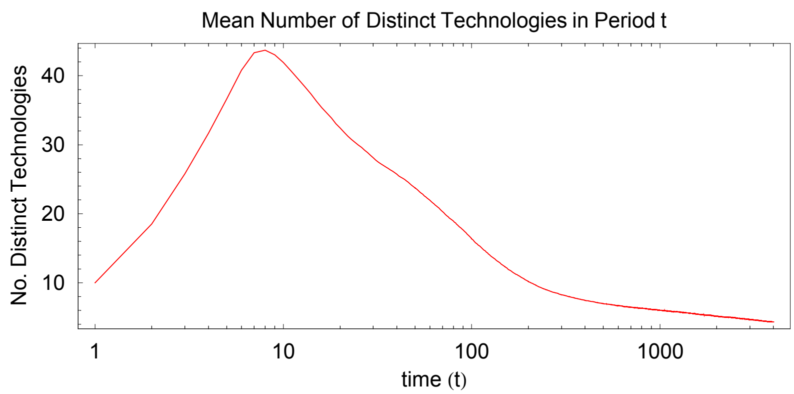

Figure 8: Number of Distinct Technologies at t = 4,000 Frequency Distribution over 1,000 Replications

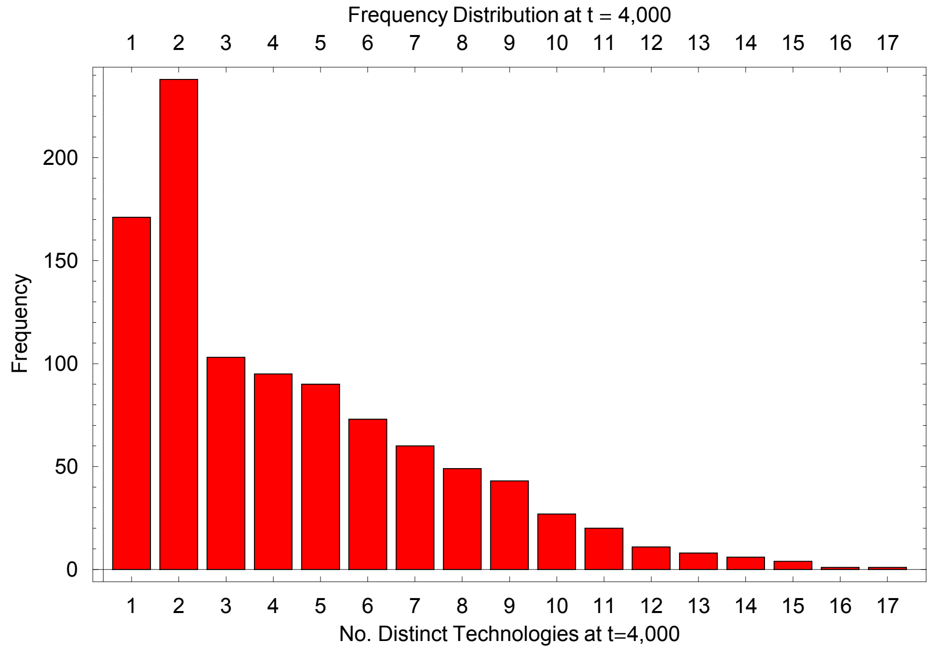

Figure 9: Mean Degree of Technological Diversity - Baseline Case

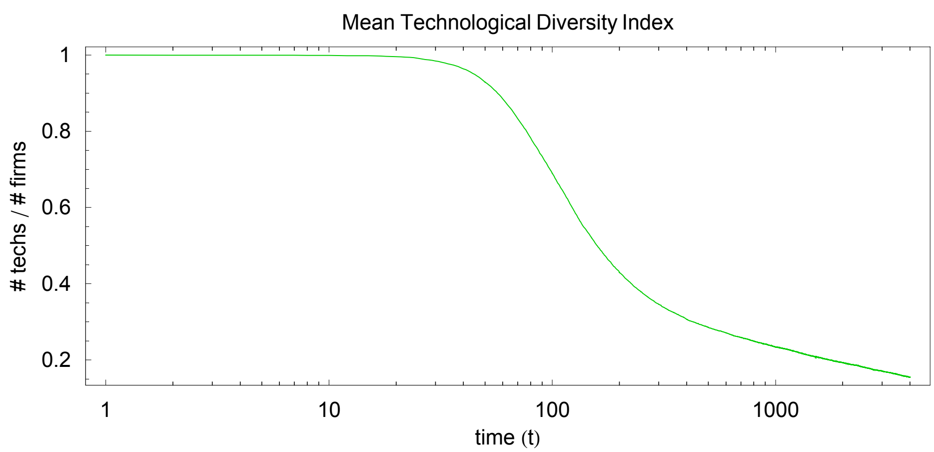

Figure 10: Impact of the Market Size (*a*)

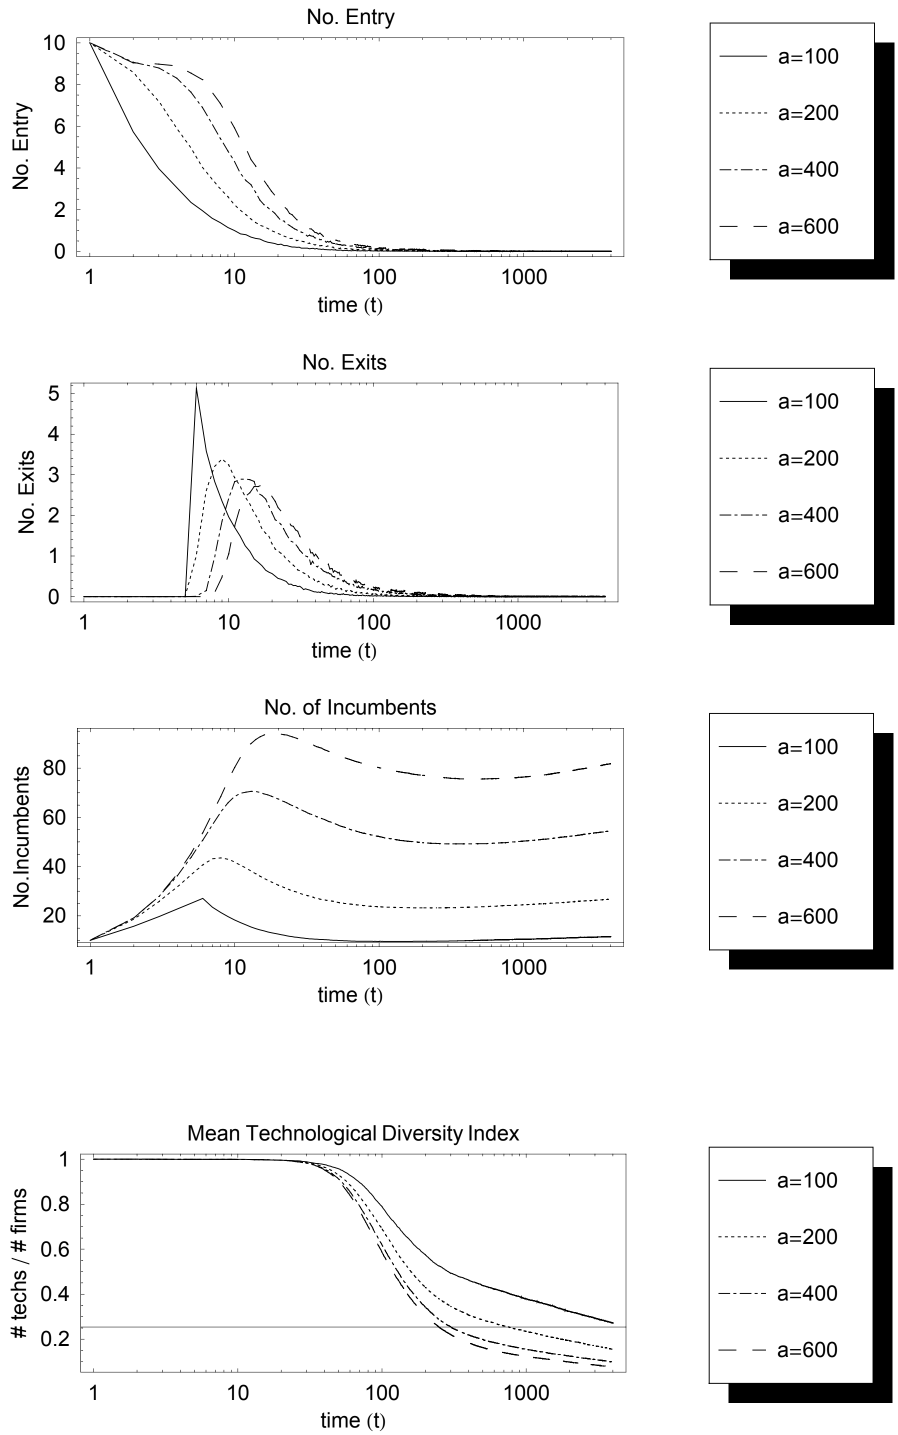

Figure 11: Impact of Fixed Cost (*f*)

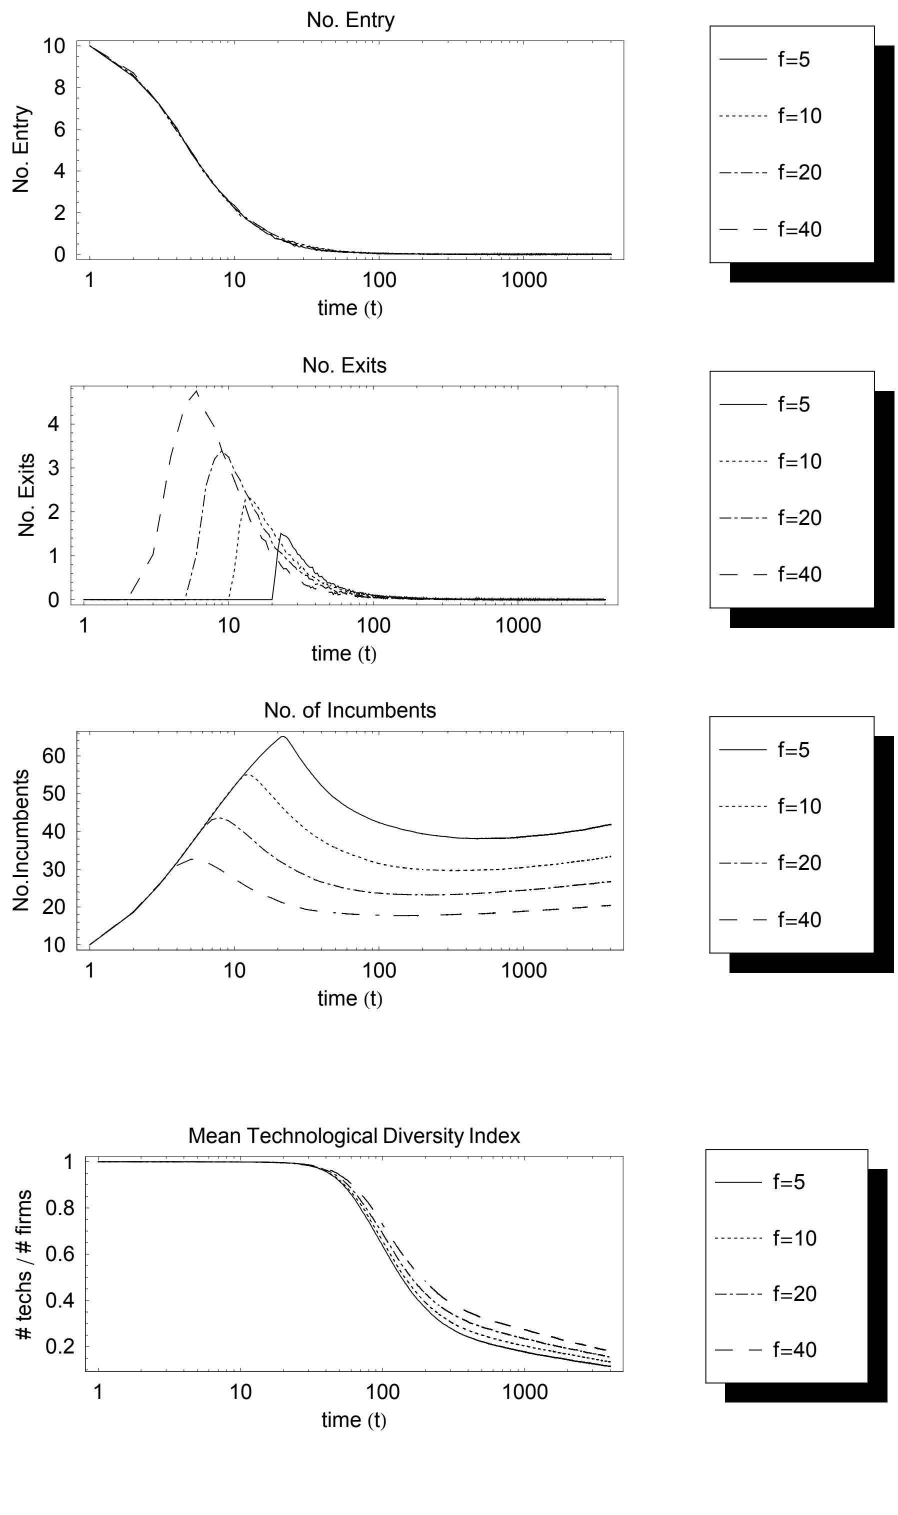

Figure 12: Impact of Potential Entrant Pool (*r*)

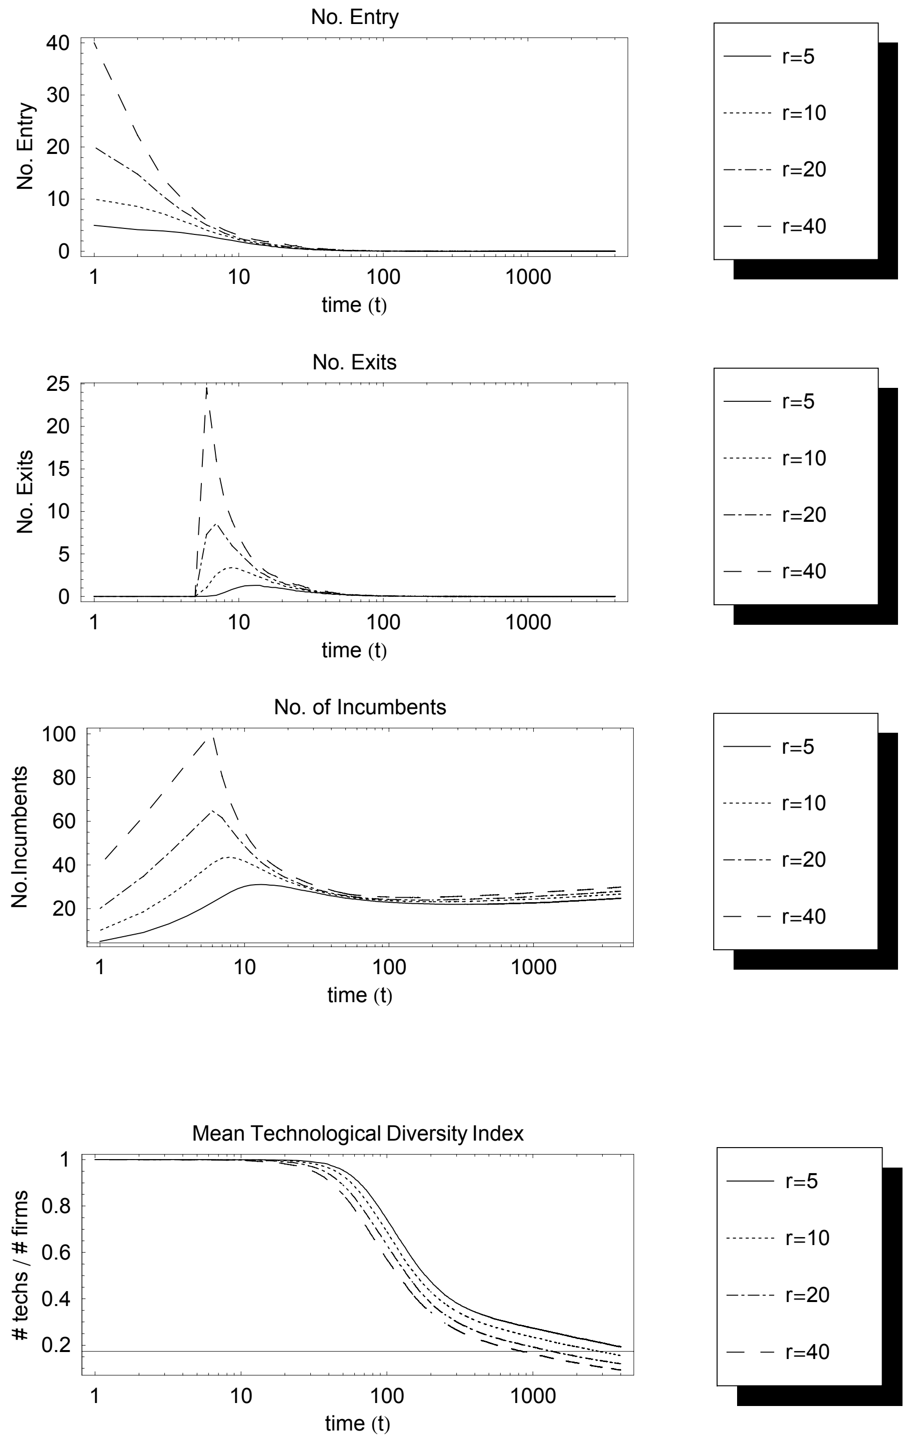

Figure 13: Impact of Start-up Fund (*b*)

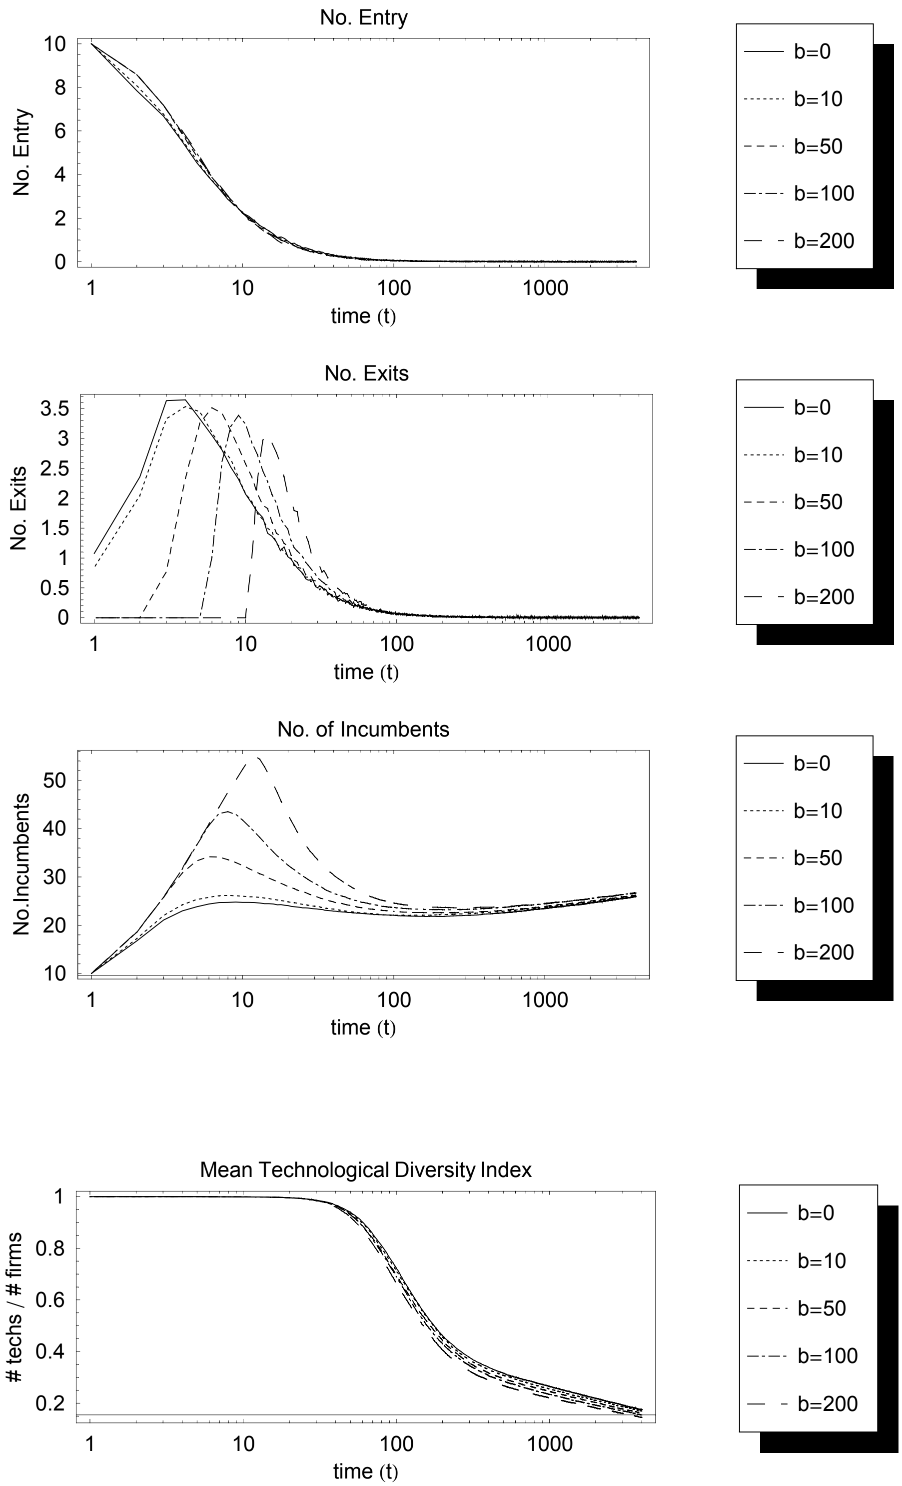

Figure 14: Impact of Search Propensity (α)

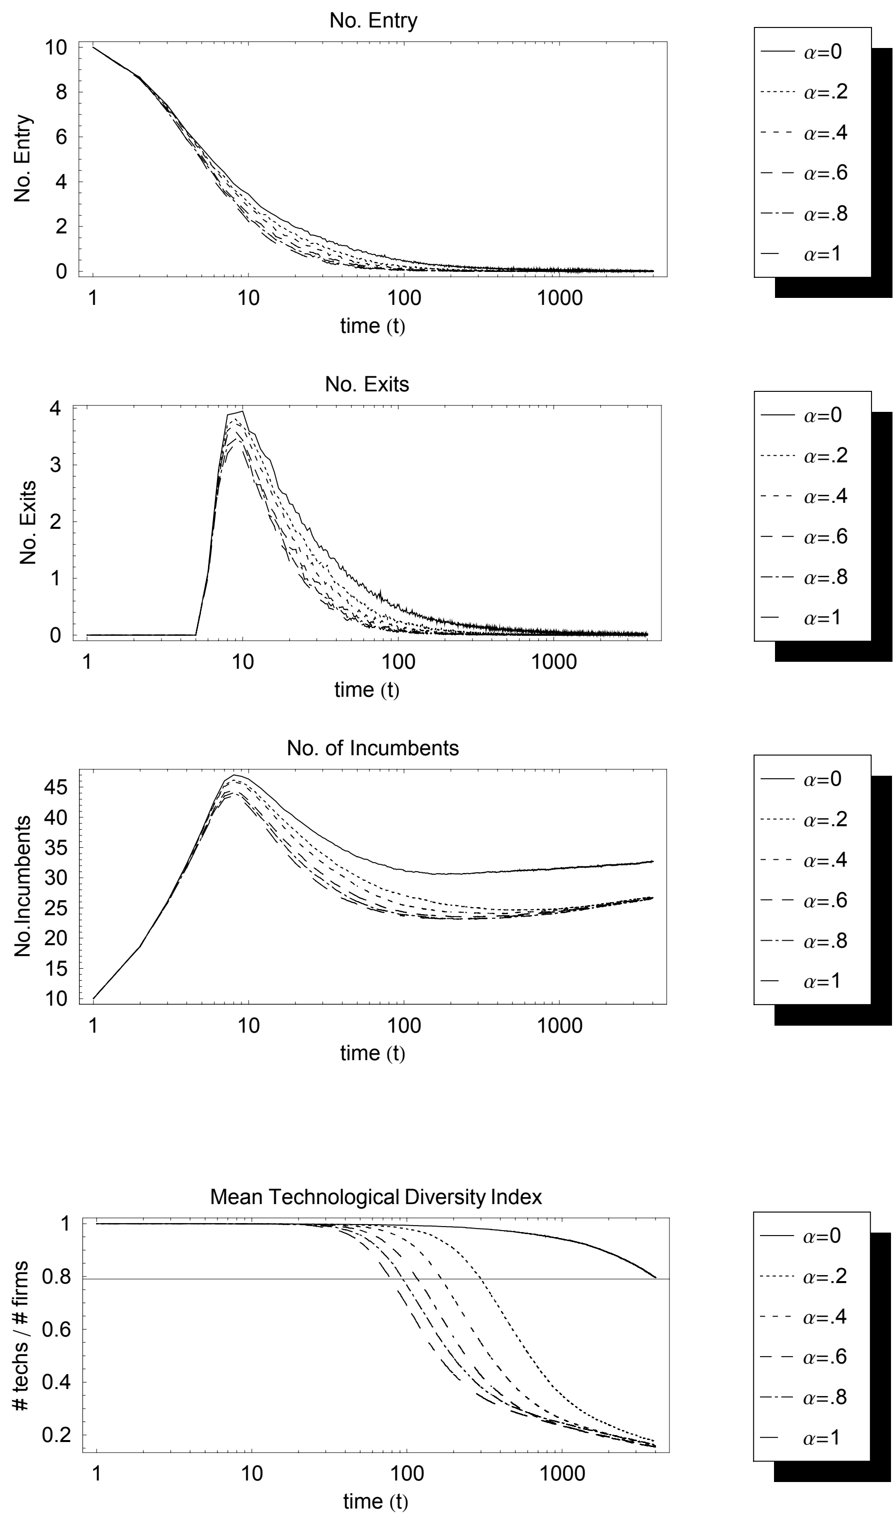

Figure 15: Impact of Search Space Size (*N*)

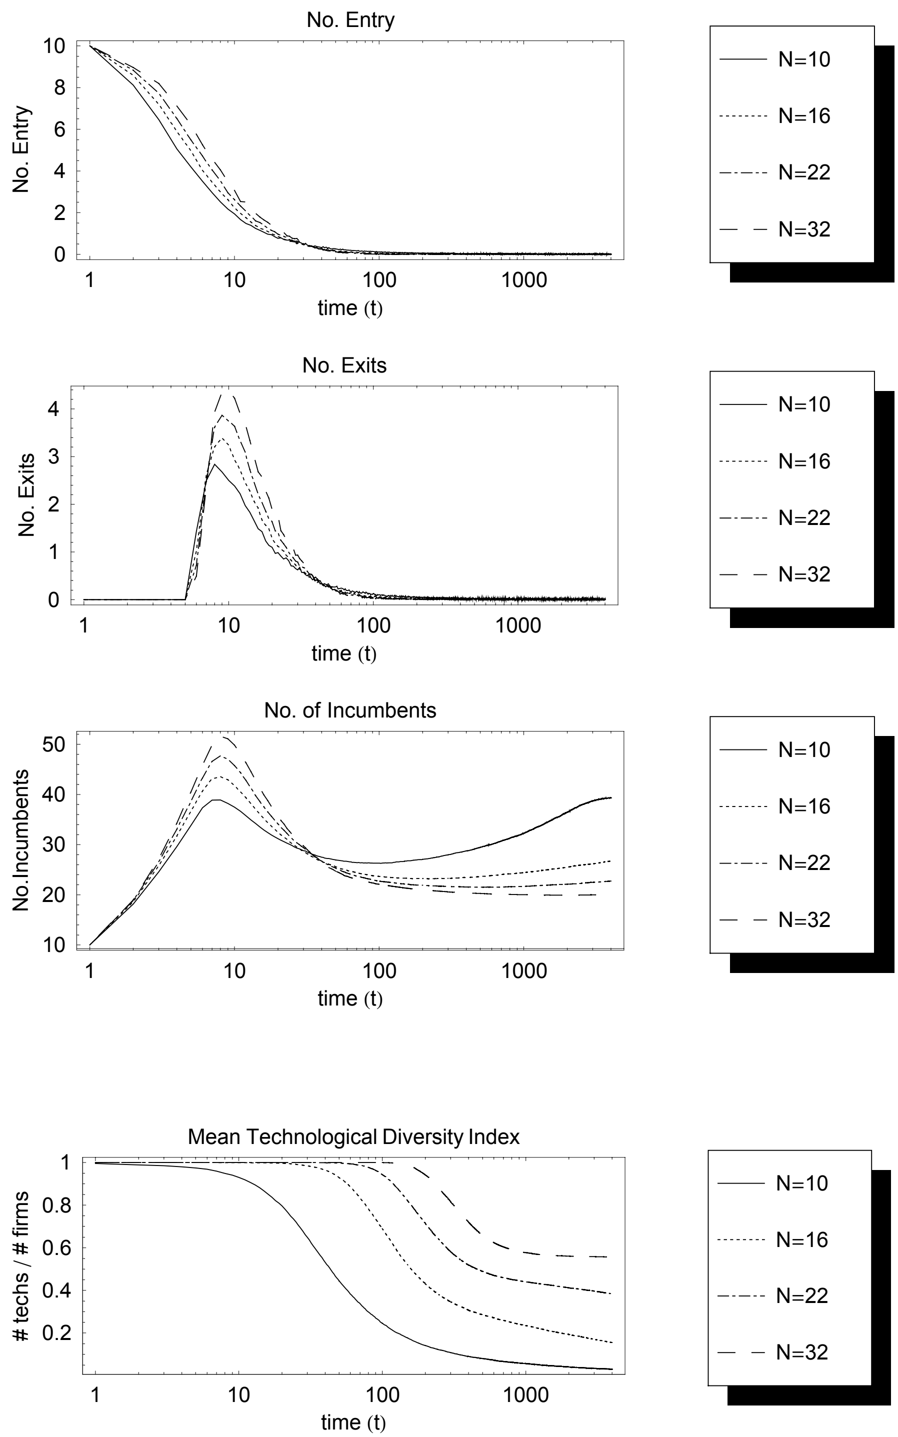

Figure 16: Impact of Technological Complexity (*K*)