J Econ Interact Coord (2009) 4:73–114
DOI 10.1007/s11403-009-0048-8

REGULAR ARTICLE

Industry dynamics with knowledge-based competition: a computational study of entry and exit patterns

Myong-Hun Chang

Received: 1 October 2008 / Accepted: 12 March 2009 / Published online: 26 March 2009
© Springer-Verlag 2009

Abstract I propose a computational model of industry evolution capable of matching many stylized facts. It views the firm as a myopic but adaptive entity whose survival depends on its ability to perform various activities with greater efficiency than its rivals. In this model, the shakeout pattern arises naturally in the early stage of industrial development. I provide a full comparative dynamics analysis of how various industry-specific factors determine the numbers and the rates of entries and exits over time as well as the ages of the exiting firms.

Keywords Industry dynamics · Shakeouts · Innovation · Imitation · Agent-based computational model

JEL Classification L10 · O30

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] A case in

[1] See Geroski (1995), Sutton (1997) and Caves (1998) for excellent surveys.

An earlier version of this paper was circulated and presented under the title, “Non-Equilibrium Industry Dynamics with Knowledge-Based Competition: An Agent-Based Computational Model.” I would like to thank the two anonymous referees, Jason Barr, Ed Bell, Jon Harford, Joe Harrington, Bill Kosteas, Edward Nell, and the conference participants at the EEA-2007 Meetings (New York, NY) as well as the 2007 International System Dynamics Conference (Boston, MA) for their helpful comments and suggestions. I gratefully acknowledge the support of the College of Graduate Studies at Cleveland State University.

M.-H. Chang (✉)
Department of Economics, Cleveland State University, Cleveland, OH 44115, USA
e-mail: m.chang@csuohio.edu
URL: http://academic.csuohio.edu/changm

point is Gort and Klepper (1982) which traces the market histories of 46 new products, empirically verifying a distinct sequence of stages 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 (a *shakeout*), 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 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,b), and Klepper (2002).

Going beyond the simple identification of the shakeout pattern, Klepper and Graddy (1990, p. 37) further noted the cross-industry differences in the entry/exit patterns over time:

> "A last observation concerns the enormous variation across new industries in the pace and severity of the prototypical pattern of industry evolution. This suggests that there are important differences across industries in the factors that condition the evolutionary process."

In a major empirical work that pulled together a set of empirical regularities involving the rates of entry and exit in a large number of U.S. manufacturing industries, Dunne et al. (1988, p. 496) also stressed the need to study how various industry-specific factors influence the evolutionary process:

> "... 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."

In this paper, I propose a model of industry evolution that is capable of generating the empirical regularities mentioned above for a wide range of parameter configurations. Specifically, it shows (1) the shakeout pattern arises naturally in the earlier part of the industry development; (2) the industry output grows at a decreasing rate, while the market price falls at a decreasing rate over time. The model also generates technological convergence among firms over the course of the industrial development; however, the extent of convergence is limited by the inherent complexity in the nature of the production system. When the entry and exit processes are compared for a wide range of industries, it is shown that: (1) the entry and exit rates over time are positively correlated; (2) a large proportion of firms that exit tends to be young. Finally, I provide a comparative dynamics analysis of how various industry-specific factors determine the numbers and the rates of entries and exits over time as well as the ages of the exiting firms.

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

The existing literature on industrial dynamics already contains works with stylized theoretical models which are capable of generating the shakeout pattern. Klepper and Graddy (1990) assumes a fixed number of potential entrants, who differ in terms of their costs and product qualities. These potential entrants are profit-maximizers with foresight 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. The shakeout pattern is then generated through the randomness in the firms' cost draws, the improvements of the cost positions through one-time 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 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.

While both Klepper and Graddy (1990) and Jovanovic and MacDonald (1994) are able to generate the shakeout pattern in a rough form, there are two aspects of their models that are not fully satisfactory. First, these models employ the standard assumption of a firm as the maximizer of expected profits with perfect foresight. Note that the type of industry environment addressed in this literature is one characterized by perpetual innovations, imperfect imitations, and constant threat of entry by outside firms. In such an environment, the incumbent firms are likely to be holding widely heterogeneous production technologies at any given point in time, and these technologies are sure to evolve over time as firms innovate and new entrants come in with new technologies of their own. Given the inherent randomness in this process, it is then difficult to accept, on conceptual grounds, the assumption of perfect rationality and foresight. Furthermore, the use of the traditional methodology with the assumption of perfect foresight severely limits the scope of their analysis: the purely analytical approach taken in their works and the consequent need to keep the analysis tractable prevent them from fully exploring the evolutionary path of the industry as well as the impacts the industry-specific factors have on the entry-exit dynamics—issues that are of significant importance to the field of industrial organization as stressed by Klepper and Graddy (1990) and Dunne et al. (1988) quoted above. Second, 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. However, in a dynamic environment, 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.

In this paper, I propose to overcome these limitations by taking a different tack to the issue of industry dynamics. The approach entails evolving an industry from its birth to maturity through a computational model, in which an ever-changing population of myopic, but adaptive, firms engage in perpetual search for improvements in their production technologies. The central feature of the model is the entry/exit dynamics that are driven by firm-level learning (via innovation and imitation) taking place in a structured search space for production technology and the market-wide competition that acts on the learning outcomes. The modeling of production technologies in this paper makes the representation of knowledge and its accumulation explicit and tractable. The rich structure of the model allows me to go beyond the simple replication of empirical regularities; it enables a systematic comparison of the entry-exit dynamics between industries which differ in their characteristics along the lines suggested by the parameter configurations considered in this paper. Consequently, I am able to respond directly to the calls made by Klepper and Graddy (1990) and Dunne et al. (1988) as quoted above. 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. Through a series of comparative dynamics exercises, I show how the industry-specific factors mentioned above affect the rates of entry and exit over time and ultimately the long-run structure of the industry.

The modeling approach taken in this paper is distinct from the more traditional approaches in two significant ways. First, it employs an agent-based computational approach with bounded rationality at the agent (firm) level. As demonstrated in Tesfatsion and Judd (2006) for a variety of research areas in economics, one of the strongest potentials of the agent-based computational modeling is that it allows a researcher to carry out detailed analyses of complex interactions among a large number of heterogeneous agents. In my model, I endow the endogenous population of heterogeneous firms with simple rules for learning (innovation and imitation), which are then 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, I make minimal demands on the level of sophistication in their reasoning abilities. 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.[3] It should be noted that my interest in this paper is *not* in characterizing the steady-state equilibrium of a mature industry, but rather in understanding the impacts that the industry-specific factors have on the very process of short- to medium-run adjustments which occur on the way to such steady-states. A similar type of exercise is carried out in Axtell (1999), in which an agent-based computational model is used to generate empirical regularities related to firm size distributions in the context of evolving firms in a given industry.

By adopting the agent-based computational approach as described above, I am dispensing with the standard notion of perfect rationality at the agent (firm) level. In its place, I adopt a set of fixed decision rules that are driven by the technologies that firms hold. An incumbent adopts a new technology (discovered through innovation or imitation) if and only if it improves the production efficiency over and above the level attained with its current technology. Likewise, a potential entrant, when contemplating entry into the market, is assumed to hold a non-negligible set of information on the relative efficiency of technologies employed in the industry and use that information in a manner that is sensible though not necessarily optimizing and certainly without "perfect foresight". This approach is in line with and is, indeed, motivated by the observation 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."

Finally, the agent-based model with these fixed decision rules enables me to work with a realistic size of agent (firm) population and to perform a rich set of comparative dynamics analyses in ways that a purely analytical (or even a computational) model with maximizing behavior and perfect foresight does not. In this regard, it is particularly useful to contrast the approach taken here to that of the Markov Perfect Equilibrium (Pakes and McGuire 1994; Ericson and Pakes 1995), another computational approach with a potential to rigorously examine the industry dynamics issue. While the MPE approach allows for forward-looking behavior on the part of the firms, it tends to suffer from the curse of dimensionality which greatly limits the number of firms that can be incorporated into the model. This limitation, of course, is a serious impediment if one is interested in producing industry dynamics that can match data. In exchange for having firms not fully forward-looking, my agent-based computational approach can allow for many firms, thereby providing a better fit to data.

The second unique aspect of the model is that the decision rules of the firms are executed in a technological environment which is constructed by the modeler *ex ante*. 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 taking place within it. More specifically, the process of producing a good is assumed to be decomposable into a system of activities, where there is

[3] 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."

a fixed number of methods available for performing each of the activities. The choice of a production technology is defined by a vector of methods, one for each activity in the system. The ultimate production efficiency of the firm using a given technology is determined by the way these methods are combined together. This 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 analyzed in theoretical literature using the concept of "supermodularity" (Milgrom and Roberts 1990, 1995; Milgrom et al. 1991).[4] However, most of the theoretical work done on this concept have focused on formally defining the conditions under which such supermodularities arise.[5] 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 presence of such complementarity within the production system is what ultimately poses a cognitive challenge for the firms as they search for ways to improve their operations, 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 system 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 technological environment within which the firms search for ways to improve their production efficiencies.[6]

The next section describes the model in detail. This is followed by a discussion in Sect. 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 Sect. 4. Section 5 offers a comparative dynamics analysis, which looks at the impacts that various classes of parameters (industry-specific factors) have on the industry dynamics identified in Sect. 4 as well as on the correlation between the rates of entry and exit within an industry. A brief summary of the results and the concluding remarks are provided in Sect. 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

[4] 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 et al. (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.

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

[6] 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).

(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 turnarounds 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.

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. I 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 $\overline{x} \equiv (x_1, \ldots, x_N)$, 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. To see the intuition behind this notion of interdependent system of activities, consider the following example of wine production as described in Porter and Siggelkow (2008):

> In the wine industry, Robert Mondavi and E. & J. Gallo compete successfully with very different systems of activities. Mondavi, the leading premium wine producer, produces high-quality wine with premium grapes, many grown in

> its own vineyards. Grapes sourced from outside growers are purchased under long-term contracts from suppliers with whom the company has deep relationships, sharing knowledge and technology extensively. Grapes are handled with great care in Mondavi's sophisticated production process, which involves extensive use of hand methods and batch technologies. Wine is fermented in redwood casks and extensively aged in small oak barrels. . ..
> Gallo, in contrast, produces large volumes of popularly priced wine using highly automated production methods. The company purchases the majority of its grapes from outside growers via arm's-length relationships and is also a major importer of bulk wine for use in blending. Gallo's production facilities look more like oil refineries than wineries. Bulk aging takes place in stainless-steel tank farms. . .. In sum, Mondavi and Gallo have chosen very different systems of contextual activities—activities that fit together and reflect the firms' different positionings.

Though the production of wine clearly involves many separate activities, the above quote focuses on a small subset of those activities that are mutually interdependent—i.e., input procurement activity, actual wine production activity, fermentation and aging activity. For each activity, the two producers have chosen different methods: Mondavi has opted for growing its own grapes (internal procurement), while Gallo has chosen to buy from outside suppliers (external procurement); Mondavi uses hand methods and batch technologies for production activity, while Gallo relies on the automated mass production technique; Mondavi for fermentation in redwood casks and aging in small oak barrels, and Gallo for bulk aging in stainless-steel tank farms. The main point is that the contribution to the overall production efficiency of a particular input procurement method depends on what methods are chosen for other interdependent activities (and vice versa). Mondavi has chosen to couple internal procurement (which ensures its tight control over the quality of grapes) with its hand methods and batch production technologies, because its efficiency gain is greater that way than when it is coupled with automated system, while Gallo has chosen to couple external procurement with an automated production system (again motivated by the efficiency considerations of its own based on their fit).[7]

In order to address the issue of fit (or complementarity) among activities in a quantitative manner, I 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.[8] Let $v_i(x_i, x_i^1, \ldots, x_i^K)$ 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, $(x_i^1, \ldots, x_i^K)$, is made explicit.[9] For expositional

[7] This example also shows how two different firms, competing in the same market, may evolve their production technologies toward two different local optima—see the discussion below of the potential for multiple local optima in the technology space.

[8] The approach employed here for incorporating the interdependence among activities in production system borrows from Kauffman (1993) NK-model, which was originally used to explore the evolutionary processes in biological systems.

[9] 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.

convenience, I will denote by $\overline{z}_i$ the vector of methods used in the activities coupled to activity i (including itself) such that $\overline{z}_i \equiv (x_i, x_i^1, \ldots, x_i^K)$.[10] The value, v_i, attached to each possible vector $\overline{z}_i$ (of length $K + 1$) is randomly chosen from [0, 100] according to a uniform distribution. The overall efficiency level of a firm, when it uses $\overline{x}$, is then

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

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*. A local optimum in this technology space is a vector, x^*, such that, for all vectors $x' \in \{\text{All } x \in X \mid D(x^*, x) = 1\}$, we have $e(x') < e(x^*)$. In other words, it is a technology for which altering the method used in any single activity will always lower the firm's efficiency level. If there is only one local optimum in the technology space, multiple firms searching independently will all eventually converge to that best technology. As has been shown by Kauffman (1993), however, the search space defined by N and K tends to have *multiple* local optima when $K > 0$.[11] Furthermore, the number of local optima, on average, increases in N and K.[12] The implication is that firms searching for efficiency improvements independently may very well end up with different technologies when $K > 0$. These properties will become relevant later on when I address the issue of technological diversity in an evolving industry. From now on, I will refer to a production system as having "a greater degree of complexity" when it has a higher value of K (and thus a greater degree of interdependence among its component activities).

Since the methods vectors evolve over time as firms successfully engage in search, I put the time superscript on them for expositional clarity: $\overline{x}_i^t \equiv (x_{i,1}^t, x_{i,2}^t, \ldots, x_{i,N}^t)$ 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 Market competition

In each period, there exists a finite number of firms that operate in the market. In this section, I define the static market equilibrium among the operating firms in any given period. For ease of presentation, I will ignore the time superscript temporarily in Sects. 2.2.1 and 2.2.2. Hence, all variables and parameters are simultaneous in this section.

[10] Given that the vector is of length $K + 1$, there are then 2^{K+1} possible vectors.

[11] The case of Mondavi vs. Gallo discussed above provides an excellent real-life example of a technology space which contains such multiple local optima.

[12] A supplement to this paper, which reports the computational results demonstrating this relationship between (N, K) and the number of local optima in the search space, is available at: http://academic.csuohio.edu/changm/main/research/papers/IDsupplement.pdf.

2.2.1 Demand and cost

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. In defining the Cournot equilibrium in this setting, I assume that all m firms produce positive quantities in equilibrium.[13] The inverse market demand function is specified to be

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

where $Q = \sum_{j=1}^{m} q_j$.

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:

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

2.2.2 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, I assume that the firms have identical fixed cost so that $f_1 = f_2 = \dots f_m \equiv f$.

[13] 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

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

$\forall i \in \{1, \ldots, m\}$, where the output vector in Cournot equilibrium is $(\widehat{q}_1, \widehat{q}_2, \ldots, \widehat{q}_m)$. 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

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

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

$\forall i \in \{1, \ldots, m\}$. 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, $\pi(\widehat{q}_i) \geq \pi(\widehat{q}_k) \forall i, k \in \{1, \ldots, m\}$.

2.3 Dynamic structure of the model

In each period of the horizon, there are four stages to the way firms make decisions. 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.

Fig. 1 Decision stages in period *t*

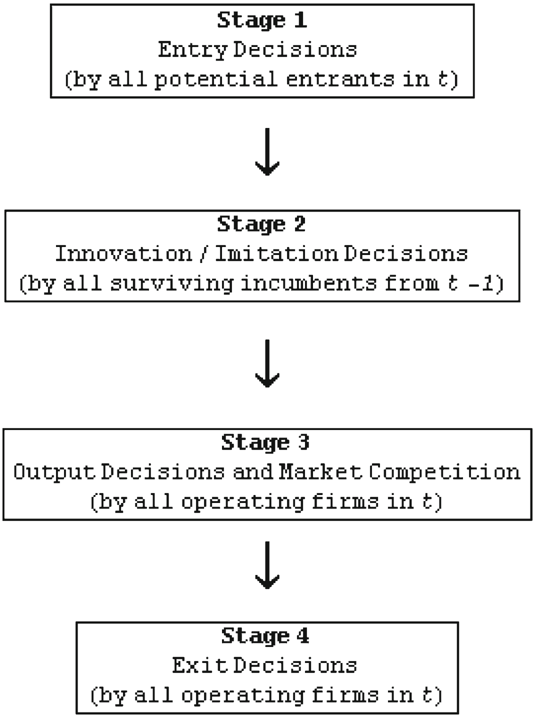

Table 1 Set notations

Notation	Definition
S^t	Set of surviving firms at the end of t
S^t_a	Those in S^t which were active in t
S^t_{-a}	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^t_+	Those in S^t which were profitable in t
L^t	Set of firms which exit the industry at the end of t

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 $S_a^{t-1} \equiv \{\text{all } j \in S^{t-1} | q_j^{t-1} > 0\}$ and $S_{-a}^{t-1} \equiv \{\text{all } j \in S^{t-1} | q_j^{t-1} = 0\}$, where $S_a^{t-1} \cup S_{-a}^{t-1} = S^{t-1}$. The inactive firms in $t-1$ are able to survive if they have sufficient wealth balance ("net capital" in my model) 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 Eq. (5).

Third, each firm $i \in S^{t-1}$ has a current "net capital" of w_i^{t-1} which it carries over from $t-1$. This net capital 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 the level of this net capital 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, I assume that the size of the potential entrants pool is fixed and constant at r throughout the entire horizon. I 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 chosen technology is its efficiency level, $e(\overline{x}_k^t)$, and the marginal cost of c_k^t for all $k \in R^t$.

2.3.1 Stage 1: entry decisions

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

$$\widehat{e}^t = \begin{cases} 0 & \text{for } t = 1, \\ \min\{e(\overline{x}_i^{t-1})\}_{\forall i \in S_a^{t-1}} & \text{for } t > 1. \end{cases} \tag{16}$$

$\widehat{e}^t$ 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:

$$\begin{cases} \text{Enter,} & \text{if and only if } e(\overline{x}_k^t) \geq \widehat{e}^t, \\ \text{Do not enter,} & \text{otherwise.} \end{cases} \tag{17}$$

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.[14] 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 Eq. (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 Eq. (5), they only need to observe the outputs and the market price to correctly infer the efficiency levels of the incumbent firms.[15]

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: $M^t = S^{t-1} \cup E^t$. I 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 capital" of b dollars. The start-up capital may be viewed as a firm's available fund that remains after paying for the one-time setup cost of entry.[16] 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 net capital 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 net capital levels for all firms in that set, $\{w_i^{t-1}\}_{\forall i \in M^t}$.

2.3.2 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, I am then capturing the exogenously specified firm propensity (common to all firms in the industry) to engage in exploration for improved technologies. This opportunity (incentive) may be

14 One could specify the entry decision rule to depend on the potential entrant's *projected* static profit based on the efficiency levels of the incumbents and his/her own, where the computation of the projected static profit is carried out *off-line*. Simulations of a similar model utilizing this entry rule also exhibited the shake-out phenomenon in exactly the same way as displayed in this paper. To avoid the excessive CPU time required for the *off-line* computations (which need to be carried out for each potential entrant for each period of the horizon), I decided to use the threshold efficiency-based entry rule as described here.

15 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, I 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.

16 The size of the one-time cost of entry is not directly relevant for my 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.

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 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 while retaining his current methods in all other activities.[17] 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_+^{t-1} denote the set of these *profitable* firms, where $S_+^{t-1} \subseteq S^{t-1}$. 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 $j \in S_+^{t-1}$ is denoted p_{ij}^t and is defined as follows:

$$p_{ij}^t \equiv \frac{\pi_j^{t-1}}{\sum_{\forall k \in S_+^{t-1}, k \neq i} \pi_k^{t-1}} \tag{18}$$

such that $\sum_{\forall j \in S_+^{t-1}, j \neq i} p_{ij}^t = 1 \forall i \in S^{t-1}$. 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:

$$\overline{x}_k^t = \begin{cases} \widetilde{x}_k^t, & \text{if and only if } e(\widetilde{x}_k^t) > e(\overline{x}_k^{t-1}), \\ \overline{x}_k^{t-1}, & \text{otherwise.} \end{cases} \tag{19}$$

Hence, a proposed technology is adopted by a firm if and only if the resulting efficiency level exceeds that of its current technology. Rather than basing the adoption decision on the expected profit maximization with perfect foresight (as in the traditional approach), the firms in my model are then assumed to be limitedly rational and rely instead on their perception of both the current and the proposed levels of production efficiency. Implicit behind this rule, however, is the assumption that the firm is able to correctly infer the efficiency level associated with the experimental technology, $\widetilde{x}_k^t$. While I believe that the firm can obtain such information only through actual experiments, I abstract away from modeling such experiments so as not to overload the model which is already complicated.

I intentionally limit the scope of technological change to one activity at a time. This implies a cognitive constraint faced by the firms.[18] As the number of activities

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

[18] These limitations in searching, both in terms of innovation and imitation, are consistent with the assumption of bounded rationality employed in this paper. The firms (agents) are viewed as being purposive—i.e., they seek improvements in their positions—but are certainly not global optimizers. Such uses of *heuristic optimization methods* have been widely observed in the recent literature on agent-based modeling—see Tesfatsion and Judd (2006) for a wide variety of examples.

in which methods can be changed rises above one, the firms are then able to consider innovations and imitations of larger scale and the interdependence among component 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.[19]

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. I adopt a version of the *Experience-Weighted Attraction* (*EWA*) learning rule as described in Camerer and Ho (1999). 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 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:

$$B_i^{IN}(t+1) = \begin{cases} \phi B_i^{IN}(t)+1, & \text{if firm } i \text{ adopted a technology through innovation in } t, \\ \phi B_i^{IN}(t), & \text{otherwise} \end{cases} \tag{20}$$

$$B_i^{IM}(t+1) = \begin{cases} \phi B_i^{IM}(t)+1, & \text{if firm } i \text{ adopted a technology through imitation in } t, \\ \phi B_i^{IM}(t), & \text{otherwise} \end{cases} \tag{21}$$

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, I 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:

$$\beta_i^t = \frac{B_i^{IN}(t)}{B_i^{IN}(t)+B_i^{IM}(t)}. \tag{22}$$

19 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, I assume that the evaluation of the technology by a firm in terms of its production efficiency is done with perfect accuracy.

The probability of pursuing imitation is, of course, $1 - \beta_i^t$. The expression in (22) implies that a favorable experience through innovation raises the probability that a firm will choose innovation again in the future.

2.3.3 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 net capital 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 the outcome is "approximated" with the static Cournot–Nash equilibrium defined in Sect. 2.2.2.

It should be noted that the direct use of the Cournot–Nash equilibrium in market competition is conceptually inconsistent with the "limited rationality" assumption employed in this paper. A more consistent approach would have been to explicitly model the process of market experimentation by myopic but adaptive firms.[20] Instead of modeling this process in microscopic detail, which would add another layer of complications, I am implicitly assuming that it is done instantly and without cost. The Cournot–Nash equilibrium is then viewed as a reasonable approximation of the outcome from that process.[21]

Recall that the equilibrium in Sect. 2.2.2 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

[20] A referee has pointed out that the modeling of the experimental behavior in output/pricing decisions would have been desirable not just for the sake of conceptual consistency, but for its capacity to face up to the empirical realities of path dependencies and lock-in effects that are frequently observed in real-world industries. I agree. The use of the Cournot–Nash equilibrium in this paper is only the first step toward exploring these possibilities. Future research will attempt to incorporate firm experimentations in both pricing and production.

[21] Note that the static Cournot equilibrium, as originally envisioned by Cournot, can be attained through a best-reply dynamic. There exists a small body of literature, in which experimental studies are conducted to determine whether firm behavior indeed converges to the Cournot–Nash equilibrium. In their pioneering work, Fouraker and Siegel (1963) conducted experiments with participants who took the role of the quantity-adjusting Cournot oligopolists under incomplete information. They did find that the Cournot–Nash equilibrium was supported in many trials for the cases of duopoly and triopoly. Similarly, Cox and Walker (1998), using linear demand and content marginal cost in Cournot duopoly, found that, if a stable equilibrium exists, then the participants in their experiments learn to play the Cournot–Nash equilibrium after only a few periods. Even though best reply dynamics do not necessarily converge in oligopolies with more than three firms (Theocharis 1960), Huck et al. (1999) finds that the best reply process does converge if firms are assumed to exhibit some *inertia* in their choice of strategy.

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 $i \in M^t$.

It should be noted that the inactive (null-production) firms have no direct impact on the short-run market outcome, since they add nothing to the industry output and the potential entrants base their entry decisions only on the production efficiency of those incumbents who were active (i.e., produced positive output) in the previous period. Nevertheless, these firms stay in the market and continue their innovation/imitation efforts *for as long as their net capital remains above the threshold level*. In the long run, this means that there is a possibility of some null-production firms making a comeback if they get a sufficiently good idea in any given period that improves their production efficiency. The null-production stage prior to exit then provides a second chance for the laggard firms, thereby having the effect of potentially reducing the severity of the shakeout relative to when the inefficient firms are forced to exit the industry immediately.

2.3.4 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' net capital levels are first updated on the basis of the profits (losses) made in t:

$$w_i^t = w_i^{t-1} + \pi_i^t. \tag{23}$$

The exit decision rule for each firm is:

$$\begin{cases} \text{Stay in} & \text{iff } w_i^t \geq d, \\ \text{Exit} & \text{otherwise,} \end{cases} \tag{24}$$

where d is the threshold level of net capital such that all firms with their current net capital 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:

$$S^t \equiv \{\text{all } i \in M^t | w_i^t \geq d\}. \tag{25}$$

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

$$L^t \equiv \{\text{all } i \in M^t | w_i^t < d\}. \tag{26}$$

The set of surviving firms, S^t, their current technologies, $\{\overline{x}_i^t\}_{\forall i \in S^t}$, and their current net capital 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 $\widehat{Q}^t \equiv \sum_{i=1}^{m^t} \widehat{q}_i^t$; (3) h^t, the Herfindahl–Hirschmann Index (HHI) in t, where $h^t \equiv \sum_{i=1}^{m^t} \left(\frac{\widehat{q}_i^t}{\widehat{Q}^t} \cdot 100 \right)^2$. 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: $\{c_i^t\}_{\forall i \in M^t}$, $\{q_i^t\}_{\forall i \in M^t}$, $\{\overline{z}_i^t\}_{\forall i \in M^t}$. 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.[22]

I run 1,000 replications for each parameter configuration, using a fresh set of random numbers for each replication.[23] I report the time paths of the above variables directly from a single typical replication as well as those averaged over 1,000 replications. Outputs are examined for a sufficiently large number of individual replications so as

Table 2 List of parameters and their values

Notation	Definition	Baseline value	Parameter values considered
N	No. of activities	16	16
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 capital for a new entrant	100	{0, 10, 50, 100, 200}
d	Threshold level of net capital for exit	0	0
α	Probability of search	1.0	{0.2, 0.4, 0.6, 0.8, 1}
T	Time horizon	2, 000	2,000

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

[23] Since I consider a horizon of 2,000 periods in each of my replications, I am then computing the outcomes from 2 million occurrences of the market competition for each parameter configuration.

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.[24]

4 The shakeout

Recall that the key empirical regularities behind the shakeout phenomenon 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 U.S. automobile tire industry, initially collected by Gort and Klepper (1982) and published in Jovanovic and MacDonald (1994), I plot in Fig. 2 the following: (a) the number of producers over 68 years from 1906 to 1973, (b) the industry output over 64 years from 1910 to 1973, and (c) the price index over 61 years from 1913 to 1973.[25] 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 Fig. 2, I now present the relevant time series outputs from a single sample run of our model, where the following configuration of parameter values are specified as the baseline: $N = 16$; $K = 2$; $r = 10$; $f = 20$; $a = 200$; $b = 100$; $d = 0.0$; $\alpha = 1.0$; $T = 2{,}000$. Hence, 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 ten 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 capital of 100. An existing incumbent leaves the market if its current net capital 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.

Figure 3 shows the time paths of the three variables—(a) number of producers, (b) industry output, and (c) 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 Fig. 2. The qualitative similarities between Figs. 2 and 3 are striking.[26] 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 Fig. 2b—the output appears to increase at

[24] 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.

[25] The complete data sets are provided in the Appendix of Jovanovic and MacDonald (1994).

[26] 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.

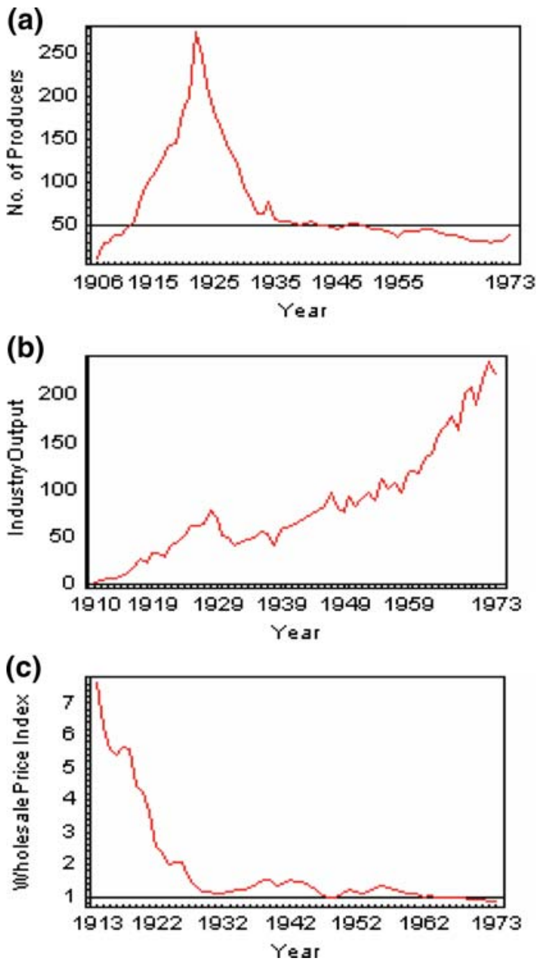

Fig. 2 Jovanovic–MacDonald data. **a** No. of producers of automobile tires, **b** Industry output for automobile tires, **c** Wholesale price index for automobile tires

a slightly increasing rate in that case. However, the simulated output path in Fig. 3b 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 the shakeout

I now present the full set of results (from multiple replications) that are generated with the baseline parameter configuration for the entire horizon of 2,000 periods. I start with the processes of entry and exit and the resulting market structure which evolves over time. Figure 4 reports the time paths over $t \in \{1, \ldots, 2{,}000\}$ of m^t, the

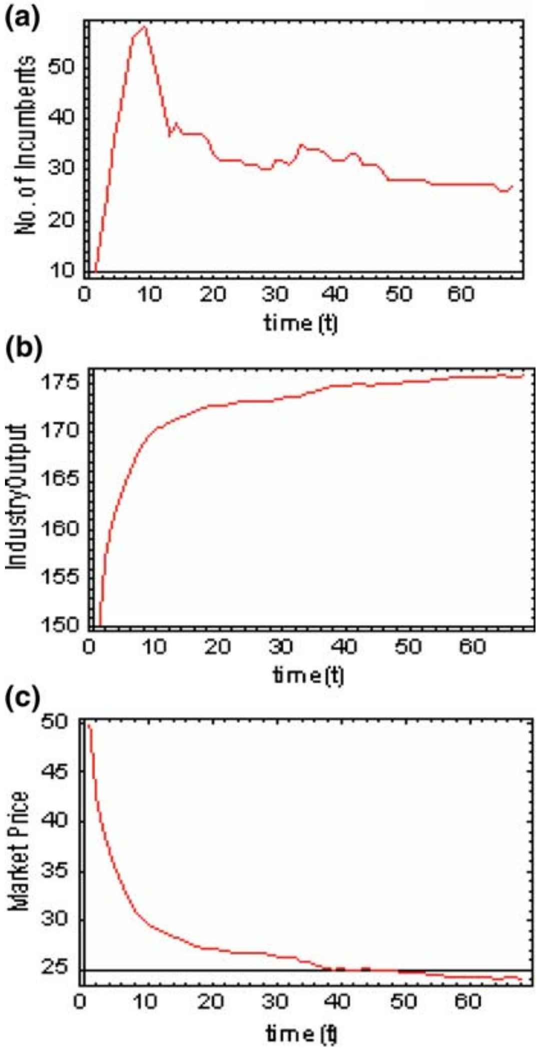

Fig. 3 Baseline results over J-M time span. **a** No. of incumbents, **b** Aggregate industry output, **c** Market price

number of incumbent firms, from 1,000 independent replications: The solid curve in the middle is the values of m^t averaged over the 1,000 replications, while the dotted curves above and below it represent the upper and lower bounds of the 95% confidence interval. Figure 5 then reports the time paths of the number of actual entrants, $|E^t|$, and the number of exits, $|L^t|$, where the values from a single typical replication are depicted in Fig. 5a, b while their averages over the 1,000 replications are shown in Fig. 5c, d. Likewise, Fig. 6 reports the market price ($\widehat{P}^t$), aggregate industry output ($\widehat{Q}^t$), and the HHI (h^t) from the same 1,000 replications with the values from a single run shown in Fig. 6a–c and the averages over 1,000 runs shown in Fig. 6d–f. All of these figures are drawn with the logarithm of time index along the horizontal axis so

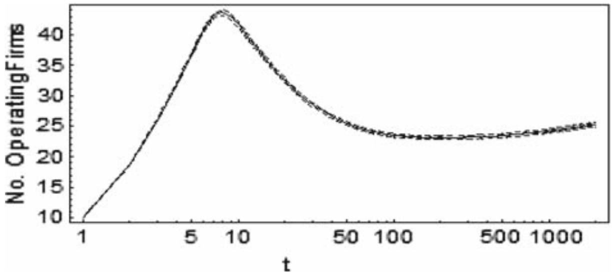

Fig. 4 Number of operating firms (baseline)

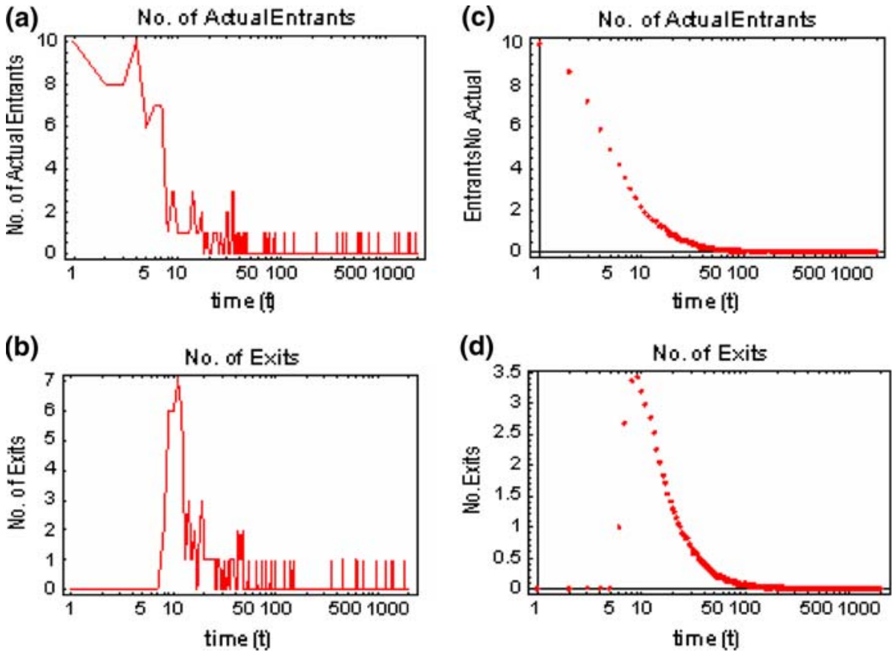

Fig. 5 Number of entrants and exits (baseline). **a** No. entrants (single run), **b** No. exits (single run), **c** No. entrants (multiple runs), **d** No. exits (multiple runs)

that the adjustments taking place in the first 100 periods are magnified relative to those in later periods.[27]

The most striking property observed in these figures is the one involving the number of incumbent firms, m^t, as shown in Fig. 4. Given that the size of the potential entrants pool is fixed at ten each period, the number of incumbent firms starts at $m^1 = 10$. Since the firms enter with a fixed start-up capital 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 mean time, 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

[27] In fact, the same scaling will be used for many of the figures presented in this paper.

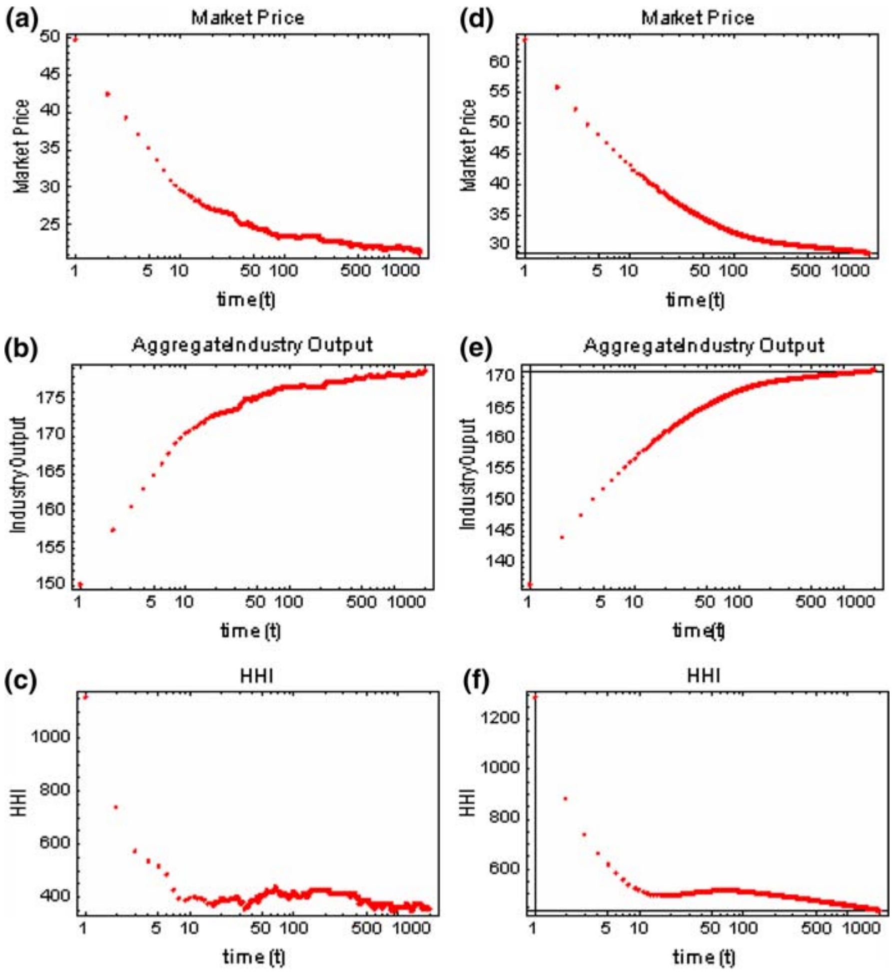

Fig. 6 Market price, industry output, and concentration. **a** Market price (single run), **b** Industry output (single run), **c** HHI (single run), **d** Market price (multiple runs), **e** Industry output (multiple runs), **f** HHI (multiple runs)

Fig. 5a, c, where the number of actual entrants, even as it declines, remains high at 5–14% of the total firms in the beginning. For the first ten 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 Fig. 5b, d, the time path of the number of exits starts to rise after the fifth period.[28] These exits of the relatively

[28] Note that the exit process does not actually start until the initial transitory phase—e.g., the first five periods in Fig. 5b, d—is over. The duration of this transitory phase (during which the inefficient incumbents are protected) is, of course, dependent upon those parameters that determine the ability of the firms to endure the economic hardship—e.g., the fixed cost which determines the size of the economic profits and losses and the size of the start-up capital which enables the new entrant to persist in the face of the continual economic losses. (This property is nicely demonstrated later on when the impacts of the fixed cost and the start-up capital on the exit time paths are visualized in Figs. 11e and 12e.)

inefficient firms occur as the general reduction in costs achieved through innovation and imitation, as well as the entry of more efficient firms, puts a continual downward pressure on the market price (see Fig. 6a, d) 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 $t = 50$. For the single replication captured in Fig. 3a the total number of incumbents at $t = 2{,}000$ is 31 firms. When averaged over 1,000 independent replications, the same measure converges to about 25 firms at $t = 2{,}000$ (Fig. 4).

It is interesting to note that while the declining price path is monotonic and smooth over the entire horizon, the exit series thus induced tends to be rather abrupt: the process of exit does not start until the initial transitory phase (the first five periods in Fig. 5b, d) is over, and then, after the initial burst, quickly declines to reach a steady-state in about 50–100 periods. The question is why the continual drop in the market price does not have the same impact on the number of exits after the initial 100 periods or so. Note that the drop in the market price is caused by a combination of two factors: (1) increase in competition by the entry of new firms with more efficient technologies; (2) improved efficiencies of the incumbent firms through innovation and imitation. While the second factor remains operative throughout the horizon (though its magnitude diminishes over time), the first factor loses its force quickly after the first 50–100 periods: As increased competition induces exits and selects more and more efficient firms to survive, it becomes tougher and tougher for potential entrants to come into the market. As shown in Fig. 5a, c, the number of entries declines significantly after the first 50–100 periods. Even though the market price still declines after this initial stage, this is mainly due to the innovation and imitation by the incumbent firms who are protected from the new entrants (with this protection becoming stronger as they continue to innovate and become more efficient). This second factor does not have the same force that the entry of new firms has on inducing the exits of the relatively inefficient firms.

While the total number of incumbents tends to decline over time after achieving the initial peak at the beginning of the industry, the aggregate industry output grows monotonically (Fig. 6b, e). This growth in output comes from the drop in the market price induced by the continual reduction in firms' marginal costs through innovation and imitation.

Property 1 *The birth of a new industry is followed by a shakeout phenomenon, in which the number of incumbent firms initially rises, then declines sharply, eventually converging to a stable level.*

Property 2 *Over the course of its development, the aggregate output of an industry grows at a decreasing rate, while the market price falls at a decreasing rate.*

The results are perfectly in line with the empirical observations made by Gort and Klepper (1982) as well as Klepper and Graddy (1990). 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 with unequal production efficiencies so that the exits of relatively inefficient firms become inevitable. Additional simulation outputs (later described in the comparative dynamics analysis) show that shakeout is a general phenomenon that arises for a wide variety of parameter configurations considered in this paper.

The time path of h^t is displayed in Fig. 6c, f. The HHI declines over time such that the market becomes increasingly less concentrated. In particular, the HHI declines steeply during the first ten periods, reflecting the rapid increase in m^t in this phase as captured in Fig. 4. The decline slows down after the initial ten periods and, then, from $t = 10$ and on it tends to remain constant. Notice from Fig. 4 that the number of incumbents, m^t, declines steeply between $t = 10$ and $t = 100$. Purely on the basis of the number of firms alone, one would expect the HHI to actually rise during these periods. My results show that this is not necessarily what happens: According to Fig. 6c, f, the time path of the HHI between $t = 10$ and $t = 100$ tends to stay relatively constant, even though the number of firms in the industry continues to drop during the same period. Why is this so? The answer lies in the degree of 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.

4.2 Marginal costs, firm outputs, and 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 $t \in \{1, \ldots, T\}$. Figure 7a captures how the distribution of marginal costs evolves over time using the data generated from a single replication (the same run that generated Fig. 3) with the baseline parameter configuration. 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, though it is not perfect (for reasons explained below). At the end of the horizon, the firms end up with 13 distinct marginal cost levels. Similarly, Fig. 7b 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 Fig. 7a.

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 Fig. 8. The time path from a single replication is reported in Fig. 8a, while the average of those from 1,000 replications is reported in Fig. 8b. 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

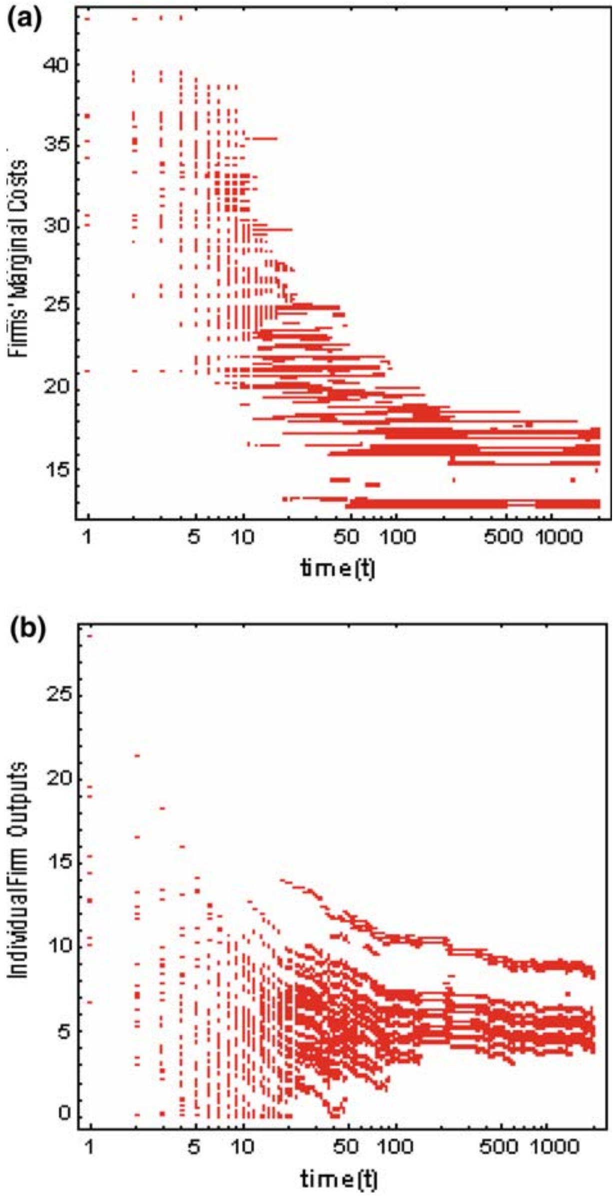

Fig. 7 Distributions of marginal costs and firm outputs. **a** Firms's marginal costs, **b** Individual firm outputs

in the number of firms 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 the corresponding time paths depicted in Figs. 4 and 8b—the number of incumbent firms stabilizes to

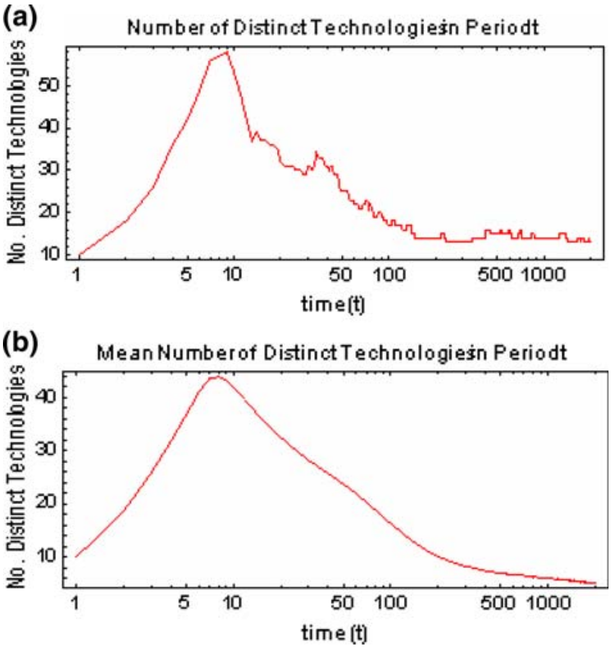

Fig. 8 Number of distinct technologies. **a** Single run, **b** Average over 1,000 runs

around 25 on average, while the number of distinct technologies declines to about 8). 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 $K > 0$. This is due to the possibility that there are multiple local optima in the search space and, hence, the technological paths of different firms may remain divergent as they evolve toward different optima over time. In the single replication captured in Fig. 8a, where $K = 2$, there are thirteen distinct technologies in existence at $t = 2{,}000$. On average, the number of distinct technologies at $t = 2{,}000$ is about 8 (Fig. 8b). It tends to have a rather wide distribution across replications. Figure 9 shows the histogram of the number of distinct technologies at $t = 2{,}000$ over 1,000 independent replications. They range anywhere from 1 to 25, although in well over 40% of the replications there are only one, two, or three distinct technologies. Diversity among firms is then not just a transient phenomenon, but is more an inherent feature of the industries that are being modeled here.

The fact that the time paths in Figs. 4 and 8b have similar shapes implies that the number of distinct technologies depends on the number of operating firms (incumbents), which changes over time. Hence, the initial sharp rise in the number of firms will automatically give rise to a proliferation of distinct technologies.[29] Since I look

[29] Although the context is a little different, it is hard to miss the strong resemblance between the initial sharp rise in the number of incumbents, accompanied by the proliferation of heterogeneous technologies, as observed in this model and the actual historical pattern in terms of *product varieties* that was observed in the U.S. automobile industry in its infancy.

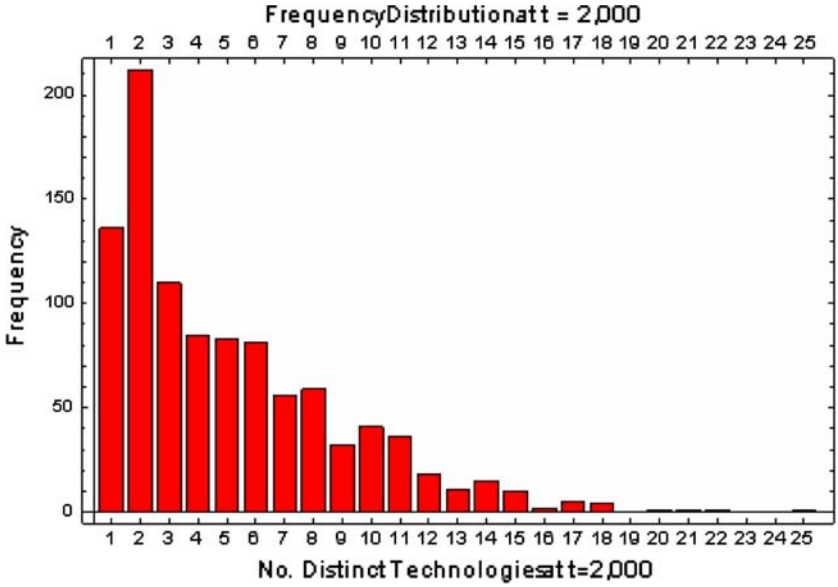

Fig. 9 Number of distinct technologies at $t = 2{,}000$ frequency distribution over 1,000 replications

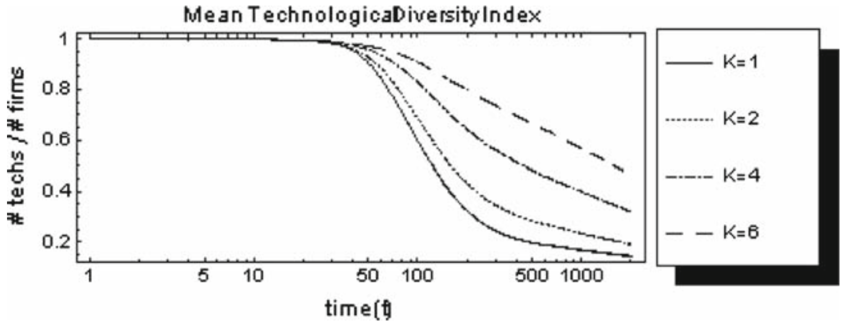

Fig. 10 Mean degree of technological diversity

for a measure of technological diversity which is consistent over time and is capable of controlling for the endogenous changes in the number of firms, I define a measure that takes the ratio of the number of distinct technologies to the number of operating firms in period t:

$$\delta(t) = \frac{\text{number of distinct technologies in } t}{\text{number of operating firms in } t \ (\equiv m^t)} \tag{27}$$

$\delta(t)$ 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 10 captures the time series of $\delta(t)$ averaged over the 1,000 replications for $K \in \{1, 2, 4, 6\}$ when all other parameters take their baseline values. The "inverted *S*" shape is clearly a general feature of $\delta(t)$.

Property 3 *From the birth of an industry to its maturity, the degree of technological diversity declines monotonically. At any given point in time, technology is more (less) diversified, when the production system has a higher (lower) degree of complexity.*

That the degree of technological diversity declines over time is not surprising. The emergence of dominant design through gradual diffusion (both in terms of product design and technology) is a well-observed phenomenon in many industries. What has not been discussed in the previous literature on industry dynamics is the second part of Property 3 which offers a theoretical insight on the relationship between the degree of complexity in technology, K, and the extent of technological diversity, $\delta(t)$. My model offers a simple intuition: Since an industry with a higher value of K (i.e., stronger inter-dependencies among activities) tends to have a larger number of locally optimal technologies in its search space, it has a higher degree of technological diversity in the long run.

5 Industry-specific factors and comparative dynamics

I now turn to a systematic investigation of how the industry-specific factors affect the entry-exit dynamics. The parameters of interest (which capture the industry-specific factors) and their values are listed in Table 2. The method is to compare the time paths of the relevant endogenous variables (e.g., numbers and rates of entry and exit, number of incumbents) by varying the value of a specific parameter, while holding the values of all other parameters constant at the baseline levels. For each given set of parameter values, I ran 1,000 replications, and averaged the time paths of endogenous variables over those replications.

5.1 The pace and severity of the shakeout

I begin by first investigating in detail how the time paths of $\left|E^t\right|$, $\left|L^t\right|$, and m^t are affected by the parameters, (a, f, r, b, α, K). The short-run impacts of some of these parameters can be inferred from the static equilibrium model of Cournot oligopoly. For instance, note from Eqs. (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 an industry are then typically inferred from the equilibrium levels of firm profitability. 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 over the course of its development? In Fig. 11a, I 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 (Fig. 11b) show that the shakeout process is

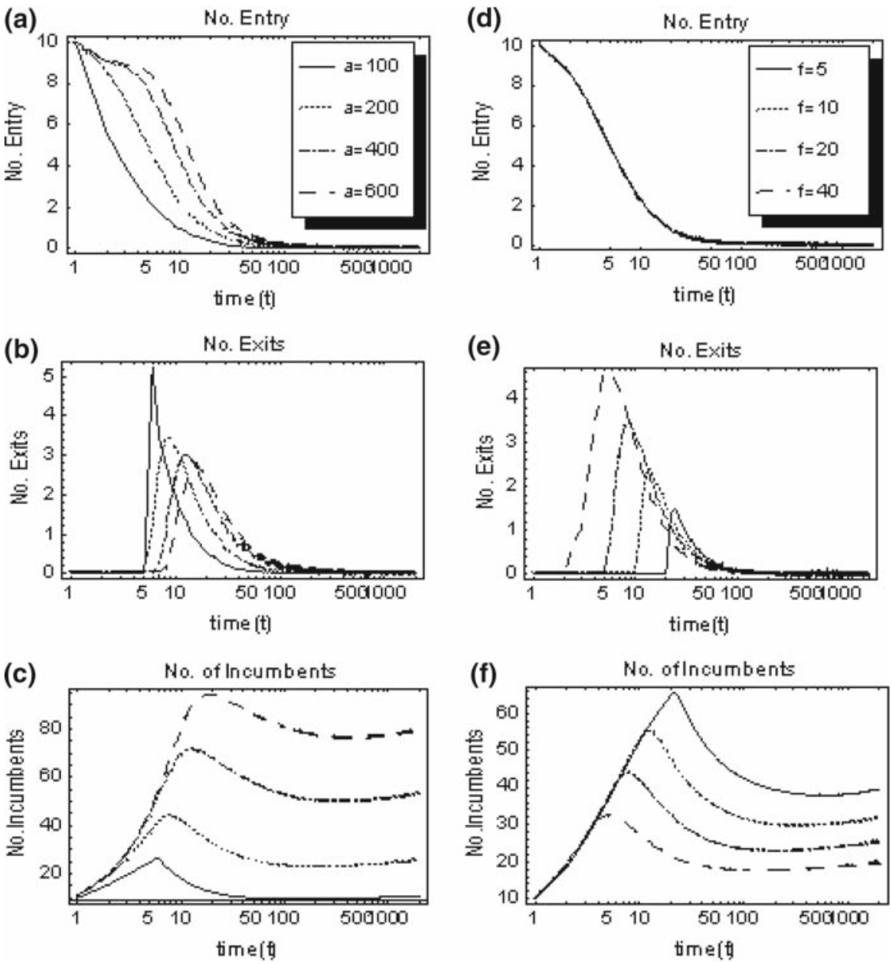

Fig. 11 Impacts of market size (a) and fixed cost (f) on the shakeout. **a** Impact of a on entry, **b** Impact of a on exit, **c** Impact of a on incumbents, **d** Impact of f on entry, **e** Impact of f on exit, **f** Impact of f on incumbents

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, do confirm this (Fig. 11c): there is a general expansion of the industry in the long run as the size of market demand increases.

Figure 11d, e, f capture the effects of 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 Fig. 11d 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 net capital). Figure 11e 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 made much more significant. Combining the entry dynamic with the exit dynamic, I then find that the number of operating firms is reduced in the long run for higher values of f (Fig. 11f): a market with a higher fixed cost tends to be more concentrated.

The impact the size of potential entrants pool (r) has on the industry dynamics is explored in Fig. 12a, b, c. 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 12a, b 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. The number of incumbent firms is significantly greater in the short run when r is higher, while that relationship, though it persists, appears to be rather weak in the long run (Fig. 12c).

It is shown in Fig. 12d that the size of the start-up capital has a negligible impact on the number of entries. However, in terms of the exit dynamics, a larger amount of start-up capital systematically delays the shakeout process by allowing the inefficient firms to linger on for a longer period of time (Fig. 12e). This, of course, increases the number of incumbents significantly in the short run, but it has a much weaker impact in the long run (Fig. 12f). Once all the exits have taken place, the resulting number of incumbents mildly increases with the size of b.

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 13a, b, c 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 lowers the number of incumbents in the short- to medium-run, though it has little impact on the long-run number of surviving firms in the market. This then implies that an industry with firms having a greater

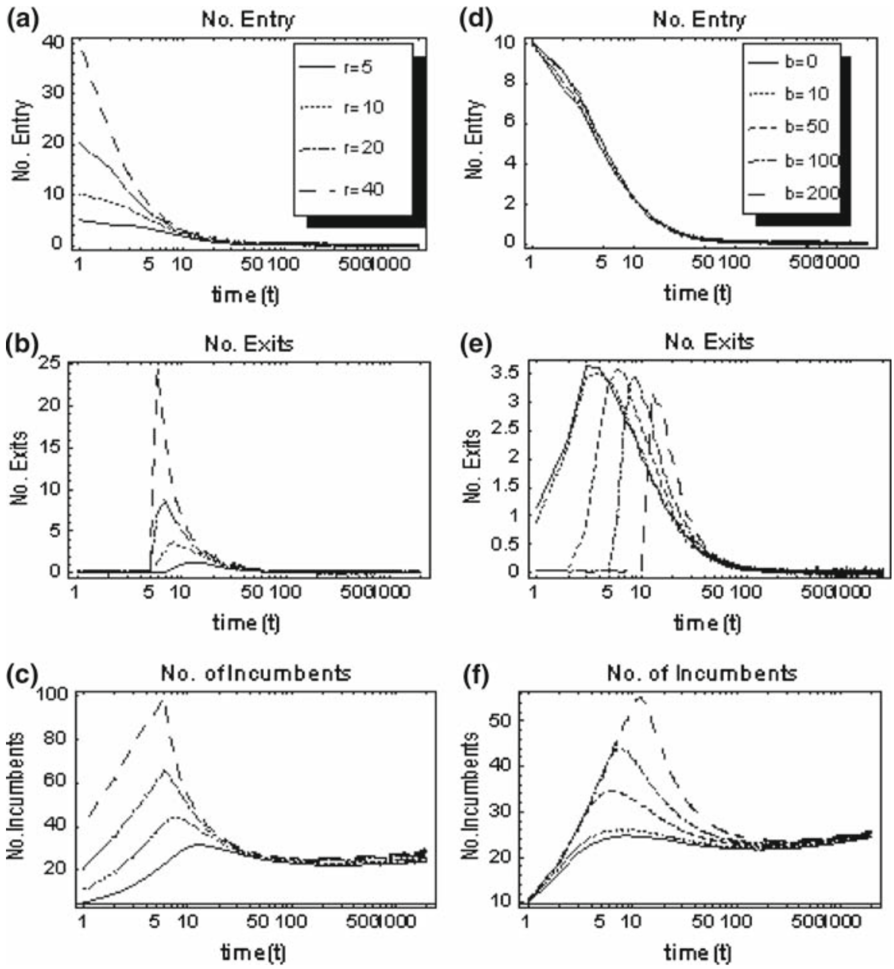

Fig. 12 Impacts of potential entrant pool (r) and start-up fund (b) on the shakeout. **a** Impact of r on entry, **b** Impact of r on exit, **c** Impact of r on incumbents, **d** Impact of b on entry, **e** Impact of b on exit, **f** Impact of b on incumbents

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.[30]

The last parameter to examine is K, the degree of interdependence among component activities of a production system. Note that an increase in K raises the number of local optima in the search space for production technologies. For any randomly chosen firm, given 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

[30] This is also directly confirmed later on when we look at the impact of α on the ages of exiting firms (Fig. 14e).

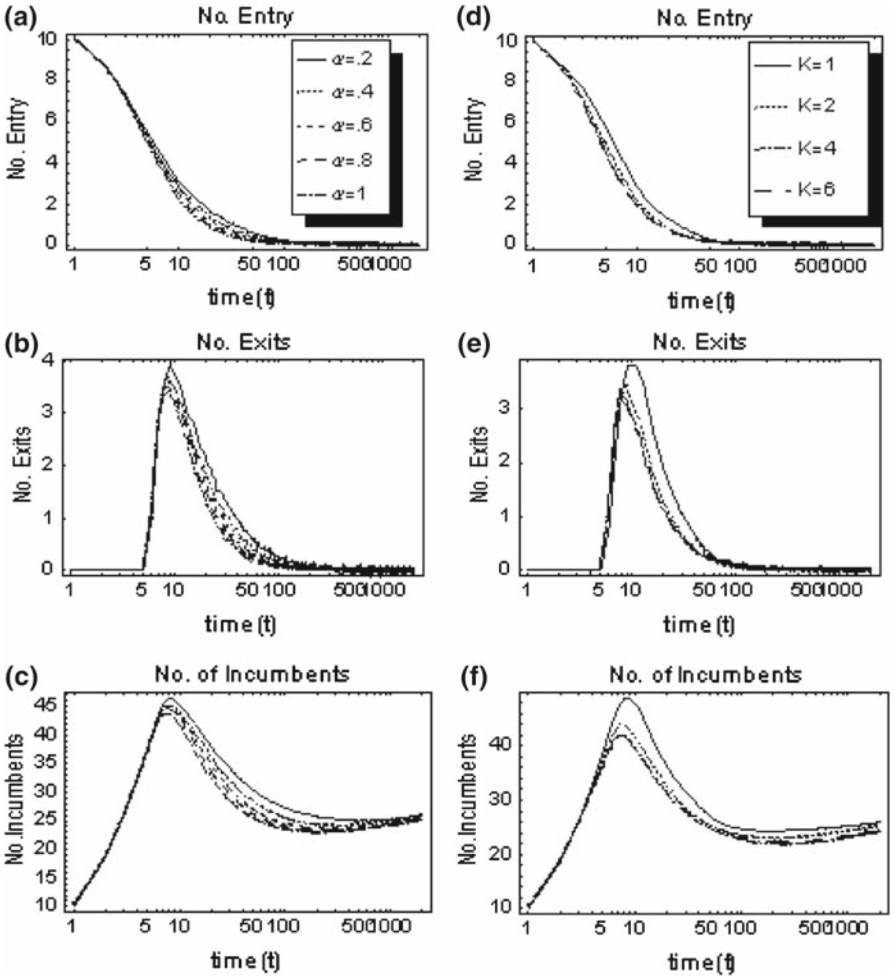

Fig. 13 Impacts of search propensity (α) and technological complexity (K) on the shakeout. **a** Impact of α on entry, **b** Impact of α on exit, **c** Impact of α on incumbents, **d** Impact of K on entry, **e** Impact of K on exit, **f** Impact of K on incumbents

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 (Fig. 13d, e). Overall, an industry then tends to be more concentrated when K is higher.

5.2 The long run market structure

Having observed the time paths of the entry and exit activities, I now focus on 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

Table 3 Total entrants, total exits, and the net entrants

Parameter values	$\sum_{t=1}^{2,000} \lvert E^t \rvert$	$\sum_{t=1}^{2,000} \lvert L^t \rvert$	$\sum_{t=1}^{2,000} \lvert E^t \rvert - \sum_{t=1}^{2,000} \lvert L^t \rvert$
a			
100	49.163 (13.4246)	38.341 (11.9041)	10.822 (2.98684)
200	93.776 (22.2481)	68.49 (19.4275)	25.286 (5.68124)
400	157.508 (39.5065)	104.395 (32.759)	53.113 (12.0115)
600	208.26 (55.4429)	128.974 (45.1026)	79.286 (17.2922)
f			
5	91.793 (23.4588)	52.294 (17.4886)	39.499 (11.1683)
10	92.837 (23.2885)	60.903 (18.7253)	31.934 (8.34118)
20	93.776 (22.2481)	68.49 (19.4275)	25.286 (5.68124)
40	90.058 (19.1232)	70.475 (17.4891)	19.583 (3.89336)
r			
5	67.517 (16.9252)	43.736 (13.912)	23.781 (5.71434)
10	93.776 (22.2481)	68.49 (19.4275)	25.286 (5.68124)
20	129.78 (31.8371)	102.673 (29.1264)	27.107 (5.82971)
40	177.433 (52.1582)	149.037 (49.5112)	28.396 (6.04779)
b			
0	99.571 (25.9273)	75.054 (22.9947)	24.517 (5.80839)
10	98.885 (24.413)	74.366 (21.496)	24.519 (5.77954)
50	96.211 (22.5263)	71.132 (19.4361)	25.079 (5.86849)
100	93.776 (22.2481)	68.49 (19.4275)	25.286 (5.68124)
200	91.137 (26.7106)	65.666 (24.3207)	25.471 (5.69698)
α			
0.2	162.201 (35.931)	136.08 (33.4396)	26.121 (5.25966)
0.4	129.281 (29.9263)	103.412 (26.986)	25.869 (5.62063)
0.6	113.314 (27.6673)	87.818 (24.7721)	25.496 (5.94957)
0.8	102.4 (23.93)	77.149 (20.9715)	25.251 (5.85613)
1.0	93.776 (22.2481)	68.49 (19.4275)	25.286 (5.68124)
K			
1	112.199 (68.6482)	86.264 (66.1245)	25.935 (5.718)
2	93.776 (22.2481)	68.49 (19.4275)	25.286 (5.68124)
4	96.714 (16.2119)	72.459 (14.0077)	24.255 (5.64817)
6	107.527 (15.5413)	83.298 (13.4685)	24.229 (5.33883)

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

occur over the entire horizon (2,000 periods)—i.e., $\sum_{t=1}^{2,000} \lvert E^t \rvert$ and $\sum_{t=1}^{2,000} \lvert L^t \rvert$. A high turnover is implied when an industry has a high total number of entries *simultaneously* with a high total number of exits. Since in my 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 I compute these values for each of the 1,000 replications, I obtain 1,000 independent realizations for each measure for a specific configuration of parameter values.

Tables 3 reports the mean values of $\sum_{t=1}^{2,000} |E^t|$, $\sum_{t=1}^{2,000} |L^t|$, and $\sum_{t=1}^{2,000} |E^t| - \sum_{t=1}^{2,000} |L^t|$ for various values of one of the parameters, (a, f, r, b, α, K), holding the values of all others constant at the baseline levels. With the exception of f and K, the total number of entries and the total number of exits move in tandem, thus implying that the industry which has a relatively high (low) total number of entries also has a relatively high (low) total number of exits over the course of its development. In addition, I make the following observations with respect to the industry-specific factors: Both the total number of entries and the total number of exits are simultaneously higher for an industry (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 capital (b) is smaller; and (4) the industry-wide search propensity (α) is lower. I also find that $\sum_{t=1}^{2,000} |L^t|$ rises monotonically in f, even though the impact of f on $\sum_{t=1}^{2,000} |E^t|$ is non-monotonic. Finally, K has a non-monotonic effect on both $\sum_{t=1}^{2,000} |E^t|$ and $\sum_{t=1}^{2,000} |L^t|$.

The industry-specific factors have monotonic impacts on the long-run number of incumbents ($\equiv \sum_{t=1}^{2,000} |E^t| - \sum_{t=1}^{2,000} |L^t|$). Property 4 summarizes these impacts. The dynamic market forces that lead to these long-run market structure were fully described in Sect. 5.1.

Property 4 *An industry sustains a greater number of firms in the long-run when*: (1) *the market demand* (a) *is larger*; (2) *the fixed cost* (f) *is lower*; (3) *the pool of potential entrants* (r) *is larger*; (4) *the start-up capital* (b) *is larger*; (5) *the industry-wide search propensity* (α) *is lower*; *and* (6) *the production system exhibits a weaker degree of interdependence* (K) *among its component activities.*

The positive relationship between the market size and the number of firms as well as the negative relationship between the fixed cost (exogenous sunk cost in Sutton 1991) and the number of firms have been observed and analyzed by a number of other works in the traditional industrial organization literature, particularly Sutton (1991) in the context of exogenous/endogenous sunk costs and the equilibrium market structure. All of these works utilize the static equilibrium models of industry competition and, hence, do not directly address the question of whether these equilibrium relationships are in fact attainable in the dynamic process of industry evolution with firm entry and exit. My model is able to explicitly *generate* these properties in a dynamic setting and, furthermore, predict how other aspects of the competition, such as the pool of potential entrants, start-up capital, search propensity, and the degree of complexity in production technologies, may affect the long-run structure of the market.

5.3 Correlation between the rates of entry and exit

In this section, I consider the *rates* of entry and exit over time, using the numbers of entries and exits as well as the number of current incumbents in each period. Let ER_k^t denote the entry rate in period t for a given replication k, where

$$ER_k^t = \frac{|E_k^t|}{m_k^{t-1}}. \tag{28}$$

Hence, ER_k^t is the number of entries occurred in period t relative to the number of firms that operated in $t-1$. Similarly, let XR_k^t denote the exit rate in period t for replication k such that

$$XR_k^t = \frac{|L_k^t|}{m_k^t}. \tag{29}$$

XR_k^t is then the number of exits relative to the total number of operating firms in that same period. Since $m_k^0 = 0$ in my model, these data are collected for $t > 1$.

Note that I performed 1,000 independent replications for each configuration of parameter values. This means that I have 1,000 separate realizations of entry and exit rates, as defined in (28) and (29), for each period from $t = 2$ to $t = 2{,}000$. For each parameter configuration, I then have 1,000 pairs of time series, $\{|E_k^t|\}_{t=2}^{2,000}$ and $\{|L_k^t|\}_{t=2}^{2,000}$. I can compute the correlation between the two time series obtained from each replication. Table 4 reports the mean and standard deviation of their correlations over the 1,000 replications. Note that the correlations are positive for all parameter configurations considered in this study. Hence, it is generally true that in those periods in which the rate of entry is high (low), the rate of exit tends to be high (low) as well. This property also implies that the peak of the entry process and the peak of the exit process are likely to be quite close to one another.[31] Note that, in my model, the entry decisions by the potential entrants are made in the first stage of a given period, while the exit decisions by the incumbent firms are made in the last stage of the same period. The positive (within-industry) correlations between the rates of entry and exit would then imply that exits tend to follow entries with no or very short lags. The intuition is, of course, that the entry by a relatively efficient firm puts an immediate downward pressure on the market price, thereby pushing the inefficient firms out of the industry.

Table 4 also shows that the magnitudes of the correlations between the two rates are systematically affected by the parameter values considered here.

Property 5 *The rates of entry and exit are positively correlated within industries. The correlation is stronger when the industry has*: (1) *a lower market demand* (a); (2) *a greater fixed cost* (f); (3) *a larger pool of potential entrants* (r); (4) *a smaller start-up capital* (b); (5) *a lower search propensity* (α); *and* (6) *a simpler production system—i.e., a weaker interdependence* (K) *among its component activities.*

These results then directly address the questions raised in Klepper and Graddy (1990) and Jovanovic and MacDonald (1994)—i.e., how do industry-specific factors determine the patterns and the rates of firm entry and exit? In particular, Property 5 shows that the conditions on the industry factors that give rise to a stronger correlation between the entry and exit rates are the ones which characterize an industry that is more prone to severe competition and generally less hospitable to new entrants.

[31] Indeed, Klepper and Simons (1997) shows this to be the case in the four U.S. industries they consider. The same has been observed by Geroski and Mazzucato (2001) using the data on the number of entrants and exitors in domestic U.S. car producers over the period of 1894–1967. They find that entry and exit are positively correlated at 0.6086. See Caves (1998, p. 1957) for additional references to works reporting similar results.

Table 4 Correlation between entry and exit rates

Parameter values	Mean correlation (SD)
a	
100	0.158083 (0.0708492)
200	0.123435 (0.0338347)
400	0.093519 (0.0224141)
600	0.0795378 (0.0186224)
f	
5	0.00719405 (0.00814931)
10	0.0418509 (0.0176444)
20	0.123435 (0.0338347)
40	0.289107 (0.0586974)
r	
5	0.0757528 (0.0289201)
10	0.123435 (0.0338347)
20	0.176745 (0.0476769)
40	0.23095 (0.0639022)
b	
0	0.705619 (0.092215)
10	0.626621 (0.108201)
50	0.250769 (0.0531736)
100	0.123435 (0.0338347)
200	0.0532195 (0.022987)
α	
0.2	0.174268 (0.0371558)
0.4	0.160763 (0.035249)
0.6	0.146683 (0.0342257)
0.8	0.134655 (0.0335614)
1.0	0.123435 (0.0338347)
K	
1	0.152033 (0.0651123)
2	0.123435 (0.0338347)
4	0.106671 (0.0276656)
6	0.0971113 (0.0259172)

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

5.4 Firm ages at the time of exit

Given that exits tend to follow entry rather quickly, one may also conjecture that those firms which exit in any given period are more likely to be recent entrants than long-surviving incumbents. We can confirm that conjecture by computing the proportion of the exiting firms (total over the entire horizon) that were of a given age or younger at the time of their exits. More specifically, I divide the maximum age of 2,000 periods into 10 equal intervals of 200 periods each. I then ask what proportion of the total exiting firms died at an age less than or equal to AGE, where AGE takes a value from 200 to 2,000 at an increment of 200. These proportions are computed for all AGEs for

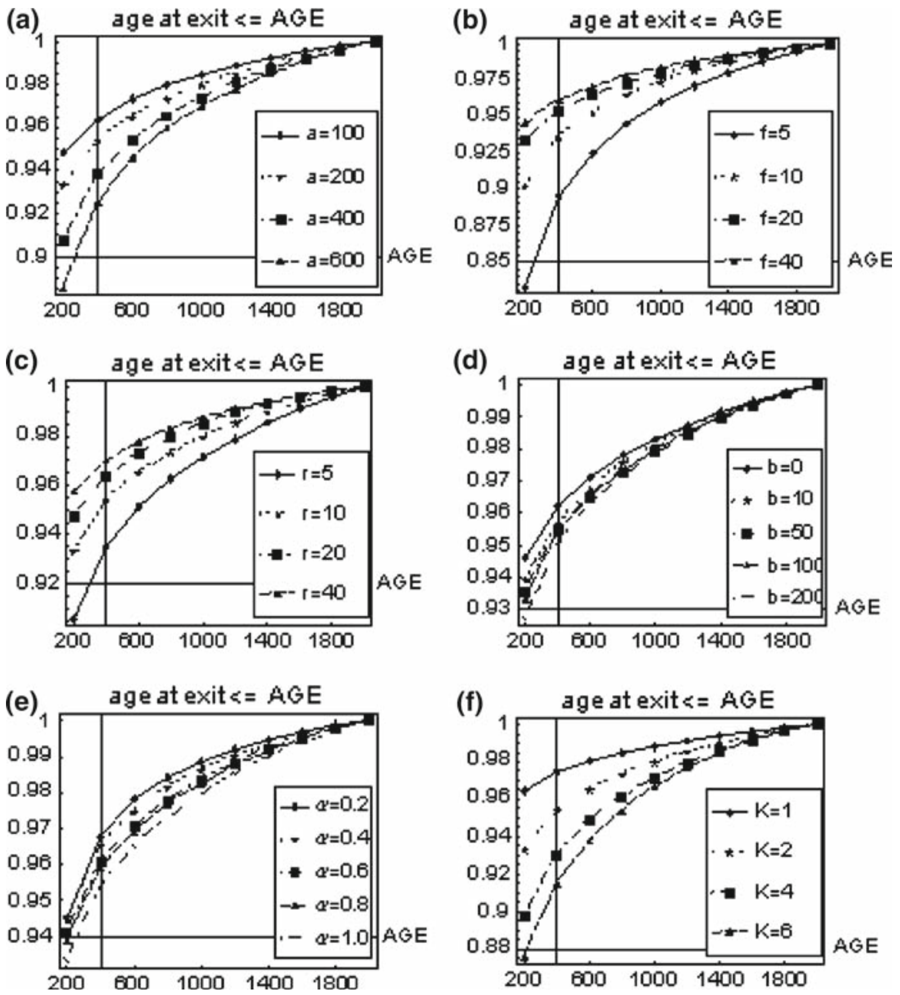

Fig. 14 Proportion of exiting firms with age ≤ AGE. **a** Impact of a, **b** Impact of f, **c** Impact of r, **d** Impact of b, **e** Impact of α, **f** Impact of K

the same parameter configurations considered in Table 4. The results are presented in Fig. 14. The first observation to be made is that, for all parameter values considered here, at least 80% of those firms which exit are of age 200 or younger.[32]

Property 6 *A large proportion of exiting firms tends to be the ones which entered the industry relatively recently.*

The following patterns can also be easily detected from these plots. For all values of AGE, the proportion of exiting firms with age less than or equal to AGE is greater when the industry has: (1) a lower market demand (a); (2) a greater fixed cost (f);

[32] The high rates of infant mortality for entrants are well-documented. See Caves (1998, pp. 1954–1959) for a brief survey.

(3) a larger pool of potential entrants (r); (4) a smaller start-up capital (b); (5) a lower search propensity (α); and (6) a simpler production system (lower K). Notice that these structural conditions are identical to those stated in Property 5: they generally signal an industry which is prone to more severe competition among its participants (and thus has a stronger correlation between the entry and exit rates).

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 a myopic but 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 technology space, 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. Of particular interest are the results concerning the turnover of the firms. I find that the shakeout pattern arises naturally in the earlier part of the industry development. Extensive comparative dynamics analyses were performed 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. Many of the industry-specific factors considered here are found to have significant impacts on the shakeout dynamics of the industry. In particular, those factors that give rise to a large number of total entries are also likely to induce a large number of total exits—an implication of which being that we should observe persistent differences in turnover rates between industries that are heterogeneous in terms of these factors. A further analysis revealed that the within-industry correlation between the entry and exit rates is positive and stronger in those industries where (1) the market demand is lower; (2) the fixed cost is higher; (3) the potential entrants pool is larger; (4) the start-up capital is smaller; (5) the firms in the industry have a lower propensity to search, and (6) the production system has a weaker interdependence among its component activities. In general, these are the conditions that represent an industry which is prone to more severe competition among its firms. It was further shown that firms tend to exit at younger ages in those industries which exhibit stronger correlations between the rates of entry and exit.

There are several ways in which this model can be enriched in a substantive manner. First, innovations and imitations are modeled 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.

Finally and more generally, rather than relying on the Cournot model as a representation of market competition and the Cournot–Nash equilibrium as the approximation of the market outcome, one can explicitly model the experimental behavior of the firms with respect to various choice variables (such as price and quantities). This would allow one to maintain the myopic but adaptive nature of the firms, fully consistent with the rest of the model, and more importantly open doors to exploring non-equilibrium behavior of the firms in the context of pricing and production decisions.

In concluding, I would like to note 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.

References

Axtell R (1999) The emergence of firms in a population of agents: local increasing returns, unstable Nash equilibria, and power law size distributions. Working Paper No. 3 (1999), Center on Social and Economic Dynamics, Brookings Institution

Bergstrom TC, Varian HR (1985) When are Nash equilibria independent of the distribution of agents' characteristics? Rev Econ Stud 52:715–718

Camerer C, Ho T-H (1999) Experience-weighted attraction learning in normal form games. Econometrica 67:827–874

Carroll GR, Hannan MT (2000) The demography of corporations and industries. Princeton University Press, Princeton

Caves RE (1998) Industrial organization and new findings on the turnover and mobility of firms. J Econ Lit XXXVI:1947–1982

Cox JC, Walker M (1998) Learning to play Cournot duopoly strategies. J Econ Behav Organ 36:141–161

Dunne T, Roberts MJ, Samuelson L (1988) Dynamic patterns of firm entry, exit, and growth. RAND J Econ 19:495–515

Epstein JM (2006) Remarks on the foundations of agent-based generative social science. In: Tesfatsion L, Judd KL (eds) Handbook of computational economics, vol 2. Agent-based computational economics handbooks in economics series. North Holland, Amsterdam

Ericson R, Pakes A (1995) Markov-perfect industry dynamics: a framework for empirical work. Rev Econ Stud 62:53–82
Ethiraj SK, Levinthal D (2004) Modularity and innovation in complex systems. Manag Sci 50:159–173
Fouraker LE, Siegel S (1963) Bargaining behavior. McGraw-Hill, New York
Geroski PA (1995) What do we know about entry? Int J Ind Organ 13:421–440
Geroski PA, Mazzucato M (2001) Modelling the dynamics of industry populations. Int J Ind Organ 19: 1003–1022
Ghemawat P (1995) Competitive advantage and internal organization: Nucor revisited. J Econ Manag Strateg 3:685–717
Gort M, Klepper S (1982) Time paths in the diffusion of product innovations. Econ J 92:630–653
Huck S, Normann H-T, Oechssler J (1999) Learning in Cournot oligopoly—an experiment. Econ J 109:C80–C95
Ichniowski C, Shaw K, Prennushi G (1997) The effects of human resource management practices on productivity: a study of steel finishing lines. Am Econ Rev 87:291–313
Jovanovic B, MacDonald GM (1994) The life cycle of a competitive industry. J Political Econ 102:322–347
Kauffman SA (1993) The origins of order. Oxford University Press, Oxford
Klepper S (2002) Firm survival and the evolution of oligopoly. RAND J Econ 33:37–61
Klepper S, Graddy E (1990) The evolution of new industries and the determinants of market structure. RAND J Econ 21:27–44
Klepper S, Simons KL (1997) Technological extinctions of industrial firms: an inquiry into their nature and causes. Ind Corp Chang 6:379–460
Klepper S, Simons KL (2000a) Dominance by birthright: entry of prior radio producers and competitive ramifications in the U.S. television receiver industry. Strateg Manag J 21:997–1016
Klepper S, Simons KL (2000b) The making of an oligopoly: firm survival and technological change in the evolution of the U.S. tire industry. J Political Econ 108:728–760
Levinthal D (1997) Adaptation on rugged landscapes. Manag Sci 43:934–950
MacDuffie JP (1995) Human resource bundles and manufacturing performance: organizational logic and flexible production systems in the World auto industry. Ind Labor Relat Rev 48:197–221
MacDuffie JP, Krafcik JF (1992) Integrating technology and human resources for high-performance manufacturing: evidence from the international auto industry. In: Kochan TA, Useem M (eds) Transforming organizations. Oxford University Press, Oxford
Milgrom P, Roberts J (1990) The economics of modern manufacturing: technology, strategy, and organization. Am Econ Rev 80:511–528
Milgrom P, Roberts J (1995) Complementarities and fit: strategy, structure, and organizational change in manufacturing. J Account Econ 19:179–208
Milgrom P, Qian Y, Roberts J (1991) Complementarities, momentum, and the evolution of modern manufacturing. Am Econ Rev 81:84–88
Pakes A, McGuire P (1994) Computing Markov-perfect Nash equilibria: numerical implications of a dynamic differentiated product model. RAND J Econ 25:555–589
Porter M (1996) What is strategy? Harv Bus Rev November–December:61–78
Porter M, Siggelkow N (2008) Contextuality within activity systems and sustainability of competitive advantage. Acad Manag Perspectives 22:34–56
Rivkin J (2000) Imitation of complex strategies. Manag Sci 46:824–844
Rivkin J, Siggelkow N (2003) Balancing search and stability: interdependencies among elements of organizational design. Manag Sci 49:290–311
Sutton J (1991) Sunk costs and market structure: price competition, advertising, and the evolution of concentration. MIT Press, Cambridge
Sutton J (1997) Gibrat's legacy. J Econ Lit XXXV:40–59
Tesfatsion L, Judd KL (eds) (2006) Handbook of computational economics, vol 2. Agent-based computational economics handbooks in economics series. North Holland, Amsterdam
Theocharis R (1960) On the stability of the Cournot solution on the oligopoly problem. Rev Econ Stud 73:133–134
Vives X (2005) Complementarities and games: new developments. J Econ Lit XLIII:437–479