Muddling through the Gales of Creative Destruction: A Non-Equilibrium Computational Model of Schumpeterian Competition*

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

September 4, 2012

Abstract

I present a computational model of R&D dynamics in an oligopoly with endogenous entry and exit. It views R&D as myopic but adaptive search for improvements in production methods and the firm as boundedly rational agent motivated by experiential learning in its pursuit of R&D. The turnover of firms, their sizes, R&D intensities, and the production efficiencies are endogenously generated and numerically tracked over time. My approach allows detailed analyses of the intra-industry relationships between the size, age, and R&D of firms as well as the inter-industry relationships between industry concentration and the R&D investments by firms.

JEL: L10

Keywords: Schumpeterian Competition, Entry, Exit, Innovation, Imitation, Market Structure

Running head: Modelling the Process of Creative Destruction

*I thank Jason Barr, Jon Harford, Matt Henry, Johannes Koenen, Robert Somogyi, and participants at the 2011 International Industrial Organization Conference (Boston), the 2012 Eastern Economic Association Conference (Boston), and the 2012 International Computational Economics & Finance Conference (Prague) for useful comments on earlier versions of this paper. I am especially grateful to Ed Bell for many stimulating conversations that proved valuable in refining the basic model. This research was funded by the Faculty Research Development Grant from the College of Graduate Studies at Cleveland State University.

> Capitalism, then, is by nature a form or method of economic change and not only never is but never can be stationary.... The fundamental impulse that sets and keeps the capitalist engine in motion comes from the new consumers' goods, the new methods of production or transportation, the new markets, the new forms of industrial organization that capitalist enterprise creates.... The opening up of new markets, foreign or domestic, and the organizational development from the craft shop and factory to such concerns as U.S. Steel illustrate the same process of industrial mutation ... that incessantly revolutionizes the economic structure *from within*, incessantly destroying the old one, incessantly creating a new one. This process of Creative Destruction is the essential fact about capitalism. — Joseph Schumpeter, *Capitalism, Socialism and Democracy*, pp. 82-83

1. Introduction

The essence of the Schumpeterian perspective is the dynamic non-equilibrium nature of the market competition; one with its focus on the *process* rather than the equilibrium the market may converge to (if ever). As such, the process of creative destruction – the persistent R&D competition combined with entry of new firms and exit of failed firms – is an on-going feature of the industrial development, rather than a mere transitional phenomenon. The research presented in this paper takes the Schumpeterian *process-view* of competition seriously and offers an integrative modeling framework within which to understand the interactive dynamics of the R&D behavior and turnover of firms over time.

The starting point for the paper is the recognition that the entries and exits by firms, as well as their R&D activities, are the permanent and persistent features of most industries. The standard equilibrium models of market competition are, however, not able to capture such non-equilibrium aspects of the competitive process. This paper represents my effort to fill this gap in the literature. In particular, I propose a computational model of industry dynamics which is capable of generating *persistent* firm R&D, as well as firm entries and exits. Using this model, I carry out a detailed study of the historical path that an industry may take from birth to maturity. The central focus of the analysis is on the endogenous intensity with which firms pursue R&D and its relationship to other endogenous variables that capture the evolving structure of the industry.

There are three parameters in the model that I focus on: 1) the size of the market demand, 2) the size of the fixed cost, and 3) the rate of change in the technological environment. These parameters are viewed as the ultimate determinants of firm behavior, having consequences for the eventual development of an industry. By observing how the industry evolves under different configurations of these parameter values, I explore the relationships between the relevant endogenous variables across industries having different characteristics.

The model has the following features: 1) The production process of a given good consists of a set of component tasks, each of which can be carried out in several different ways; 2) Firms can differ in terms of how they produce a given good – i.e., they may produce the same good using different combinations of methods; 3) The purpose of R&D is to reduce the cost of producing the good by finding a better method for carrying out each component task of the production process; 4) The production environment is subject to random but permanent changes caused by factors external to the industry – i.e., a method which used to work well for a given task may no longer be adequate and vice versa. The central driver of the Schumpeterian dynamics, in this framework, is *the stochastic nature of the technological environment*. The mechanism works as follows. In any given period, a random technological change can differentially affect the effectiveness of the "methods" used by firms for various production tasks. Given the heterogeneity in the firms' technologies, some firms may benefit from the change, while others may lose from it. It also creates an opportunity for potential entrants to come in with a unique set of production methods well-suited for the new environment, while simultaneously providing the incumbent firms with renewed opportunities to search for improvements in production methods by engaging in R&D. What follows is an episode of industrial *shakeout* that entails a wave of firm entries and exits. When these exogenous technological shifts occur at a constant rate, the industry experiences *waves* of shakeouts with "persistent" series of entries and exits as well as that of R&D investments by surviving incumbents. My goal is to identify and explain any patterns that may exist in this process.

As I study the computationally-generated histories of the industries, I ask the following questions: 1) Within an industry, what are the evolving relationships between firm sizes, firm ages, and their R&D intensities? 2) Across industries, how do different demand and cost characteristics affect the rate of turnover, the R&D intensity, and the industry structure? 3) What are the long-run relationships between these endogenous variables? How are these relationships determined by the exogenous rate of change in the technological environment? 4) When there are two distinct modes of R&D (innovative R&D vs. imitative R&D), how does a firm's choice between them evolve over time, and how is it affected by the industry-specific factors?

The computational model predicts that younger firms engage in more R&D than the established older firms, and are also more efficient in general. It also predicts that the firm size, as measured by the sales revenue, is positively correlated with the firm's R&D intensity. In terms of the inter-industry comparisons, the model first predicts that the industries with smaller market sizes and/or larger fixed costs experience greater rate of firm turnover and are more concentrated on average. The market size and the fixed cost also affect the endogenous R&D activities in the same way: The average R&D expenditure per firm, which captures the intensity of R&D, is larger in industries having smaller market sizes and/or larger fixed costs. As to the choice between the different modes of R&D, the model predicts that the ratio of the innovation expenditure to the imitation expenditure is higher in industries with smaller market size and/or larger fixed cost. Collectively, these findings suggest that the industries with greater turbulence – i.e., higher rate of firm turnover – are more concentrated, and firms in such industries are more active in R&D; in addition, these firms display a bias toward innovative R&D over imitative R&D in the long run.

This paper is organized as follows. A brief review of the background literature is given in the next section. Section 3 presents the formal model. The design of computational experiments performed in the paper is described in Section 4. Section 5 presents the results. How intra-industry asymmetries in firm size, age, and efficiency relate to the endogenous R&D intensity is discussed in Section 5.1. Section 5.2 provides a comparative analysis of the relationship between industrial structure and the R&D intensity, when industries are differentiated by the size of the market demand and the fixed cost. The robustness of the results and how they are affected by the degree of turbulence in the technological environment are briefly discussed in Section 5.3. Section 6 offers concluding remarks.

2. Background Literature

A large body of empirical literature on R&D has accumulated. The central driver of this literature has been the two hypotheses of Schumpeter: 1) the intensity of a firm's R&D is positively related to the firm's size; 2) the intensity of the aggregate R&D and the degree of industry concentration are positively related. Cohen and Levin (1989), in a comprehensive survey of this literature, concludes that the relationship between firm size and R&D is inconclusive, while the relationship between concentration and R&D is generally positive and significant. The underlying models for these empirical studies have been the static equilibrium models. Given the process of *creative destruction* as envisioned by Schumpeter, both firm size and industry structure should be viewed as coevolving with the R&D activities of the firms over the course of the industrial development as the firms enter and exit the industry as part of that process. The observed empirical regularities may then be viewed as snapshots taken along such a developmental path. Computationally evolving an industry, the proposed model offers a petri-dish approach to studying the factors that generate such regularities.

Separately from the literature on R&D, there exists a body of empirical studies that explore the turnover of firms in various industries. This literature has identified many regularities that are at odds with the predictions of many equilibrium-based models. For instance, Caves (2007, p.9) points out the widely observed positive correlations between the contemporary rates of entry and exit. He notes: "Turnover in particular affects entrants, who face high hazard rates in their infancy that drop over time. It is largely because of high infant mortality that rates of entry and exit from industries are positively correlated (compare the obvious theoretical model that implies either entry or exit should occur but not both). The positive entry-exit correlation appears in cross-sections of industries, and even in time series for individual industries, if their life-cycle stages are controlled." Furthermore, the variations in the extent of these regularities appear to imply that the industry-specific factors play a major role in the way a given industry develops. In their study of the shakeout patterns in new industries, Klepper and Graddy (1990, p.37) notes: "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." Dunne *et al.* (1988, p.496) expresses a similar sentiment as they study the regularities involving the rates of entry and exit in a large number of U.S. manufacturing industries: "... 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." By incorporating entry, exit, and R&D into a setting in which technological environment can fluctuate over time, the proposed model can generate an industry dynamic that is consistent with these observations.

The modeling approach taken in this paper is an agent-based computational one in which firms are assumed to be boundedly rational and driven by a set of fixed decision rules. As shown in Tesfatsion and Judd (2006), 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 the proposed model, an endogenous population of heterogeneous firms is endowed with simple rules for adaptive learning (R&D), combined with the rules governing the entry and exit behaviors of the potential entrants and the incumbent firms. 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.

By adopting the agent-based computational approach, I am dispensing with the standard notion of perfect rationality at the agent (firm) level. An alternative approach would have been to use a Markov-Perfect Equilibrium (MPE) model [Pakes and McGuire (1994); Ericson and Pakes (1995)], in which firms are endowed with full rationality and perfect foresight. While this approach can, in theory, address the similar issues explored in this project, it suffers from the well-known "curse of dimensionality," greatly limiting the number of firms and the number of parameters that can be incorporated into the model [Doraszelski and Pakes (2007)].[1] This limitation is a serious impediment if one is interested in producing industry dynamics that can match data. In exchange for having firms be not fully forward-looking, the agent-based computational approach taken in this paper will allow for many firms, thereby providing a better fit to data.

A closely related work to this paper is Nelson and Winter (1982) on the evolutionary theory of firms. As in this paper, they take a computational approach and focus on the non-equilibrium process of the market competition in the presence of firm R&D. The boundedly rational nature of the firm's decision-making is also recognized as being central to the process of innovation and market competition. However, they do not explicitly study the process of entry and exit by firms. Rather, once an initial number of firms is specified, any change in the industry structure is inferred from a change in the Herfindahl-Hirschmann Index based on the evolving capital stocks of the fixed incumbent population. By specifying a fixed number of firms as an initial condition, their work then precludes the possibility of studying the transient state of a new infant industry or the intertemporal patterns in the firm turnovers. In this paper, I intend to investigate the development path of an industry from its birth to maturity with the process of entry and exit made fully explicit.

[1] For efforts to overcome this difficulty within the MPE framework, see Weintraub *et al.* (2008, 2010).

3. A Model of Industry Dynamics

The base model entails an evolving population of firms which interact with one another through repeated market competition. Central to this process are the heterogeneous production technologies held by the firms and the R&D mechanism through which they evolve over time. This section describes the unique way in which production technology is modelled, as well as the multi-stage market process within which R&D decisions are made.

3.1. Basic Features

3.1.1. Technology

In each period, firms engage in market competition by producing and selling a homogeneous good. The good is produced through a process that consists of N distinct tasks. Each task can be completed using one of two different methods. Even though all firms produce a homogeneous good, they may do so using different combinations of methods for the N component tasks. The method chosen by the firm for a given task is represented by a bit (0 or 1) such that there are two possible methods available for each task and thus 2^N variants of the production technology. In period t, a firm's *technology* is then fully characterized by a binary vector of N dimensions which captures the complete set of methods it uses to produce the good. Denote it by $\underline{z}_i^t \in \{0,1\}^N$, where $\underline{z}_i^t \equiv (z_i^t(1), z_i^t(2), ..., z_i^t(N))$ and $z_i^t(h) \in \{0,1\}$ is firm i's chosen method in task h.

In measuring the degree of heterogeneity between two technologies (i.e., method vectors), $\underline{z}_i$ and $\underline{z}_j$, we use "Hamming Distance," which is the number of positions for which the corresponding bits differ:

$$D(\underline{z}_i, \underline{z}_j) \equiv \sum_{h=1}^{N} |z_i(h) - z_j(h)| . \tag{3.1}$$

The crucial perspective taken in this model is that the efficiency of a given technology depends on the environment it operates in. In order to represent the technological environment that prevails in period t, I specify a unique methods vector, $\underline{\widehat{z}}^t \in \{0,1\}^N$, which is defined as the *optimal technology* for the industry in t. How well a firm's chosen technology performs in the current environment depends on how close it is to the prevailing optimal technology in the technology space. More specifically, the marginal cost of firm i realized in period t is specified to be a direct function of $D(\underline{z}_i^t, \widehat{\underline{z}}^t)$, the Hamming distance between the firm's chosen technology, $\underline{z}_i^t$, and the optimal technology, $\widehat{\underline{z}}^t$. The firms are uninformed about $\widehat{\underline{z}}^t$ *ex ante*, but engage in search to get as close to it as possible by observing their marginal costs. The optimal technology is common for all firms – i.e., all firms in a given industry face the same technological environment. As such, once it is defined for a given industry, its technological environment is completely specified for all firms since the efficiency of any technology is well-defined as a function of its distance to this optimal technology.

I allow turbulence in the technological environment. Such turbulence is assumed to be caused by factors external to the industry in question such as technological innovations that originate from outside the given industry.[2] The external technology shocks *redefine* firms' production environment and such environmental shifts affect the cost positions of the firms in the competitive marketplace by changing the effectiveness of the methods they use in various activities within the production process. These unexpected disruptions then pose renewed challenges for the firms in their efforts to adapt and survive. It is precisely this kind of external shocks that I try to capture in this paper. My approach is to allow the optimal technology, $\widehat{\underline{z}}^t$, to vary from one period to the next, where the frequency and the magnitude of its movement represent the degree of turbulence in the technological environment. The exact mechanism through which this is implemented is described in Section 3.2.1.

Finally, in any given period t, the optimal technology is unique. While the possibility of multiple optimal technologies is a potentially interesting issue, it is not explored here because in a turbulent environment, where the optimal technology is constantly changing, it is likely to be of negligible importance.[3]

[2] In a framework closer to the neoclassical production theory, one could view an externally generated innovation as a shock that affects the relative input prices for the firms. If firms, at any given point in time, are using heterogeneous production processes with varying mix of inputs, such a change in input prices will have diverse impact on the relative efficiencies of firms' production processes – some may benefit from the shock; some may not. Such an external shock will then require (with varying degrees of urgency) a series of adaptive moves by the affected firms for their survival.

[3] Chang (2009) offers an alternative approach by modeling the technological environment as being stable but with

3.1.2. Demand, Cost, and Competition

In each period, there exists a finite number of firms that operate in the market. In this subsection, I define the static market equilibrium among such operating firms. The static market equilibrium defined here is then used to *approximate* the outcome of market competition in each period. In this sub-section, I temporarily abstract away from the time superscript for ease of exposition.

Let m be the number of firms operating in the market. The firms are Cournot oligopolists, who choose production quantities of a homogeneous good. In defining the Cournot equilibrium in this setting, I assume that all m firms produce positive quantities in equilibrium.[4] The inverse market demand function is:

$$P(Q) = a - \frac{Q}{s}, \tag{3.2}$$

where $Q = \sum_{j=1}^{m} q_j$ and s denotes the size of the market.[5]

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

$$C(q_i) = f_i + c_i \cdot q_i. \tag{3.3}$$

For simplicity, the firms are assumed to have identical fixed cost: $f_1 = f_2 = ... = f_m = f$.

The firm's marginal cost, c_i, depends on how different its technology, $\underline{z}_i$, is from the optimal technology, $\widehat{\underline{z}}$. Specifically, c_i is defined as follows:

$$c_i(\underline{z}_i, \widehat{\underline{z}}) = 100 \cdot \frac{D(\underline{z}_i, \widehat{\underline{z}})}{N}. \tag{3.4}$$

Hence, c_i increases in the Hamming distance between the firm's chosen technology and the optimal technology for the industry. It is at its minimum of zero when $\underline{z}_i = \widehat{\underline{z}}$ and at its maximum of 100

multiple locally optimal technologies. The main focus is on the industry dynamics during the initial shakeout phase, where one of the objectives was to investigate the impact of multiple optima on the shakeout dynamics. In the current paper, I am more interested in the dynamics of R&D and firm turnover along the steady-state path in the presence of technological turbulence. As such, I abstract away from the possibility of multiple local optima.

[4] This assumption is made strictly for ease of exposition in this section. 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 Section 3.2.2

[5] This function can be inverted to $Q = s(a - P)$. For a given market price, doubling the market size then doubles the quantity demanded.

when all N bits in the two technologies are different from one another. The total cost can then be re-written as:

$$C(q_i) = f + 100 \cdot \frac{D(\underline{z}_i, \widehat{\underline{z}})}{N} \cdot q_i. \tag{3.5}$$

Given the demand and cost functions, firm i's profit is:

$$\pi_i(q_i, Q - q_i) = \left(a - \frac{1}{s} \sum_{j=1}^{m} q_j \right) \cdot q_i - f - c_i \cdot q_i. \tag{3.6}$$

Taking the first-order condition for each i and summing over m firms, we derive the equilibrium industry output rate, which gives us the equilibrium market price, $\overline{P}$, through equation (3.2):

$$\overline{P} = \left(\frac{1}{m+1} \right) \left(a + \sum_{j=1}^{m} c_j \right). \tag{3.7}$$

Given the vector of marginal costs defined by the firms' chosen technologies and the optimal technology, $\overline{P}$ is uniquely determined and is independent of the market size, s. Furthermore, the equilibrium market price depends only on the *sum* of the marginal costs and not on the *distribution* of c_is.

The equilibrium firm output rate is:

$$\overline{q}_i = s \left[\left(\frac{1}{m+1} \right) \left(a + \sum_{j=1}^{m} c_j \right) - c_i \right]. \tag{3.8}$$

Note that $\overline{q}_i = s \left[\overline{P} - c_i \right]$: A firm's equilibrium output rate depends on its own marginal cost and the equilibrium market price. Finally, the Cournot equilibrium firm profit is

$$\pi(\overline{q}_i) = \overline{P} \cdot \overline{q}_i - f - c_i \cdot \overline{q}_i = \frac{1}{s} \left(\overline{q}_i \right)^2 - f \tag{3.9}$$

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

3.2. Dynamic Structure of the Model

In the beginning of any typical period t, the industry opens with two groups of decision makers: 1) a group of incumbent firms surviving from $t-1$, each of whom enters t with a technology, $\underline{z}_i^{t-1}$, and its net wealth, w_i^{t-1}, carried over from $t-1$; 2) a group of potential entrants ready to consider entering the industry in t, each with an endowed technology of $\underline{z}_j^t$ and its start-up wealth. All firms face a common technological environment within which his/her technology will be used. This environment is fully represented by the prevailing optimal technology, $\widehat{\underline{z}}^t$, which is not necessarily the same as $\widehat{\underline{z}}^{t-1}$.

Central to the model is the view that the firms engage in search for the optimal technology over time, but with limited foresight. What makes this "perennial" search non-trivial is the stochastic nature of the production environment – that is, the technology which was optimal in one period is not necessarily optimal in the next period. This is captured by allowing the optimal technology, $\widehat{\underline{z}}^t$, to vary from one period to the next in a systematic manner. The mechanism that guides this shift dynamic is described next.

3.2.1. Turbulence in the Technological Environment

Consider a binary vector, $\underline{x} \in \{0,1\}^N$. Define $\delta(\underline{x}, l) \subset \{0,1\}^N$ as the set of points that are exactly Hamming distance l from $\underline{x}$. The set of points that are *within* Hamming distance l of $\underline{x}$ is then defined as

$$\Delta(\underline{x}, l) \equiv \bigcup_{i=0}^{l} \delta(\underline{x}, i). \tag{3.10}$$

The following rule governs the shift dynamic of the optimal technology:

$$\widehat{\underline{z}}^t = \begin{cases} \widehat{\underline{z}}' & \text{with probability } \gamma \\ \widehat{\underline{z}}^{t-1} & \text{with probability } 1-\gamma \end{cases}$$

where $\widehat{\underline{z}}' \in \Delta(\widehat{\underline{z}}^{t-1}, g)$ and γ and g are constant over all t.[6] Hence, with probability γ the optimal technology shifts to a new one within g Hamming distance from the current technology, while with

[6]For the computational experiments reported in this paper, $\widehat{\underline{z}}'$ is chosen from $\Delta(\widehat{\underline{z}}^{t-1}, g)$ according to uniform distribution.

probability $1-\gamma$ it remains unchanged at $\underline{\widehat{z}}^{t-1}$. The volatility of the technological environment is then captured by γ and g, where γ is the rate and g is the maximum magnitude of changes in technological environment.

The change in technological environment is assumed to take place in the beginning of each period before firms make any decisions. While the firms do not know what the optimal technology is for the new environment, they are assumed to get accurate signals of their *own* marginal costs based on the new environment when making their decisions to enter or to perform R&D. This is clearly a strong assumption. A preferred approach would have been to explicitly model the process of learning about the new technological environment; it is for analytical simplicity that I abstract away from this process.

3.2.2. Multi-Stage Decision Structure

Each period consists of four decision stages – see Figure 1. Denote by S^{t-1} the set of surviving firms from $t-1$, 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. The inactive firms in $t-1$ survive to t if and only if they have sufficient net wealth to cover their fixed costs in $t-1$. Each firm $i \in S^{t-1}$ possesses a production technology, $\underline{z}_i^{t-1}$, carried over from $t-1$, which gave rise to its marginal cost of c_i^{t-1} as defined in equation (3.4). It also has a current net wealth of w_i^{t-1} it carries over from $t-1$.

Let R^t denote a finite set of *potential* entrants who contemplate entering the industry in the beginning of t. 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, $\underline{z}_k^t$, randomly chosen from $\{0,1\}^N$ according to uniform distribution. In addition, each potential entrant has a fixed start-up wealth it enters the market with.

The definitions of the set notations introduced in this section and used throughout the paper are summarized in Table 1.

Stage 1: Entry Decisions In stage 1 of each period, the potential entrants in R^t first make their decisions to enter. Just as each firm in S^{t-1} has its current net wealth of w_i^{t-1}, we will let $w_j^{t-1} = b$ for all $j \in R^t$ where b is the fixed "start-up" wealth common to all potential entrants. The start-up wealth, b, may be viewed as a firm's available fund that remains after paying for the one-time set-up cost of entry.[7] 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 negative value.

It is important to specify what a potential entrant knows as it makes the entry decision. A potential entrant k knows its own marginal cost, c_k^t, based on the new environment, $\widehat{z}^t$.[8] It also has observations on the market price and the incumbent firms' outputs from $t-1$ – that is, P^{t-1} and $q_i^{t-1} \forall i \in S^{t-1}$. Given these observations and the fact that $\overline{q}_i = s[\overline{P} - c_i]$ from equation (3.8), k can infer c_i^{t-1} for all $i \in S^{t-1}$. While the surviving incumbent's marginal cost in t may be different from that in $t-1$, I assume that the potential entrant takes c_i^{t-1} to stay fixed for lack of information on $\widehat{z}^t$. The potential entrant k then uses c_k^t and $\{c_i^{t-1}\}_{\forall i \in S^{t-1}}$ in computing the post-entry profit expected in t.

Given the above information, the entry rule for a potential entrant takes the simple form that it will be attracted to enter the industry if and only if it perceives its post-entry net wealth in period t to be strictly positive. The entry decision then depends on the profit that it expects to earn in t following entry, which is assumed to be the static Cournot equilibrium profit based on the marginal costs of the active firms from $t-1$ and itself as the only new entrant in the market.[9]

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

[8] It is not that the potential entrant k knows the content of $\widehat{z}^t$ (the optimal method for each activity), but only that it gets an accurate signal on c_k^t (which is determined by $\widehat{z}^t$).

[9] That each potential entrant assumes itself to be the only firm to enter is clearly a strong assumption. Nevertheless, this assumption is made for two reasons. First, it has the virtue of simplicity. Second, Camerer and Lovallo (1999) provides some support for this assumption by showing in an experimental setting of business entry that most subjects who enter tend to do so with overconfidence and excessive optimism.

The decision rule of a potential entrant $k \in R^t$ is then:

$$\begin{cases} \text{Enter,} & \text{if and only if } \pi_k^e(\underline{z}_k^t) + b > \underline{W}; \\ \text{Stay out,} & \text{otherwise,} \end{cases} \tag{3.11}$$

where π_k^e is the static Cournot equilibrium profit the entrant *expects* to make in the period of its entry and $\underline{W}$ is the threshold level of wealth for a firm's survival (common to all firms).

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 which will guarantee some threshold level of profits given its belief about the market structure and the technological environment. Denote by M^t the set of firms ready to compete in the industry: $M^t \equiv S^{t-1} \cup E^t$. At the end of stage 1 of period t, we then have a well-defined set of competing firms, M^t, with their current net wealth, $\{w_i^{t-1}\}_{\forall i \in M^t}$ and their technologies, $\underline{z}_i^{t-1}$ for all $i \in S^{t-1}$ and $\underline{z}_j^t$ for all $j \in E^t$.

Stage 2: R&D Decisions In stage 2, the surviving incumbents from $t-1$, S^{t-1}, engage in R&D to improve the efficiency of their existing technologies. Given that the entrants in E^t entered with new technologies, they do not engage in R&D in t. In addition, only those firms with sufficient wealth to cover the R&D expenditure engage in R&D. I will denote by I_i^t the R&D expenditure incurred by firm i in t.

The R&D process transforms the incumbent's technology from $\underline{z}_i^{t-1}$ to $\underline{z}_i^t$, where $\underline{z}_i^t = \underline{z}_i^{t-1}$ if either no R&D is performed in t or R&D is performed but its outcome is not adopted. The modeling of this transformation process is described separately in Section 3.4 for expositional clarity.

Stage 3: Output Decisions and Market Competition Given the R&D decisions made in stage 2 by the firms in S^{t-1}, all firms in M^t now have the updated technologies $\{\underline{z}_i^t\}_{\forall i \in M^t}$ as well as their current net wealth $\{w_i^{t-1}\}_{\forall i \in M^t}$. With the updated technologies, the firms engage in Cournot competition in the market, where we "approximate" the outcome with the Cournot equilibrium defined in Section 3.1.2.[10]

Note that the equilibrium in Section 3.1.2 was defined for m firms under the assumption that all m firms produce positive quantities. In actuality, given the asymmetric costs, there is no reason to think that all firms in M^t will produce positive quantities in equilibrium. Some relatively inefficient firms may shut down their plants and stay inactive. What we need is 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 done in the following sequence of steps. Starting from the initial set of active firms, compute the equilibrium outputs for each firm. If the outputs for one or more firms are negative, then de-activate the least efficient firm from the set of currently active firms – i.e., set $q_i^t = 0$ where i is the least efficient firm. Re-define the set of active firms (as the previous set of active firms minus the de-activated firms) and recompute the equilibrium outputs. Repeat the procedure until all active firms are producing non-negative outputs. Each *inactive* firm produces zero output and incurs the economic loss equivalent to its fixed cost. Each *active* firm produces its equilibrium output and earns the corresponding profit. We then have π_i^t for all $i \in M^t$.

Stage 4: Exit Decisions Given the single-period profits or losses made in stage 3 of the game, the firms in M^t consider exiting the industry in the final stage. Each firm's net wealth is first updated on the basis of the profits (or losses) made in stage 3 as well as the R&D expenditure incurred in stage 2:[11]

$$w_i^t = w_i^{t-1} + \pi_i^t - I_i^t, \tag{3.12}$$

[10] I admit to the use of Cournot-Nash equilibrium as being conceptually inconsistent with the "limited rationality" assumption employed in this paper. However, explicitly modeling the process of market experimentation would further complicate an already complex model. As such, I implicitly assume that experimentation is done instantly and without cost. Cournot-Nash equilibrium is then assumed to be a reasonable approximation of the outcome from that process.

[11] It does not matter whether the R&D expenditure is subtracted from the net wealth in stage 2 or in stage 4. It is a sunk cost by the time market competition starts and, as such, it has no impact on the firm's output decision in stage 3.

where I_i^t is the firm's R&D expenditure. The exit decision rule for each firm is then:

$$\begin{cases} \text{Stay in} & \text{iff } w_i^t \geq \underline{W}, \\ \text{Exit} & \text{otherwise,} \end{cases} \tag{3.13}$$

where $\underline{W}$ is the previously-defined threshold level of net wealth such that all firms with their current net wealth below $\underline{W}$ exit the market. Define L^t as the set of firms which exit the market in t. Once the exit decisions are made by all firms in M^t, the set of surviving firms from period t is then defined as:

$$S^t \equiv \{\text{all } i \in M^t | w_i^t \geq \underline{W}\}. \tag{3.14}$$

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

3.3. Prior Work Using the Base Model

Variants of the base model, as described above, have been used in two earlier papers, Chang (2009, 2011). In both papers, the second stage R&D activity was assumed to be *exogenous* and *costless*. More specifically, R&D was viewed as serendipitous discovery, in which the methods used in one or more of the tasks were randomly altered for experimentation each period. Chang (2009) assumed a stable technological environment in which the optimal technology did not change from one period to next – i.e., $\gamma = 0$. Instead, the technology itself was assumed to be *complex* in nature such that there were multiple optima. The main focus was on the *shakeout* phase of an industry's life cycle, immediately following its birth. The model generated shakeout patterns consistent with the empirical observations.

Chang (2011) allowed turbulence in technological environment as in this paper – i.e., $\gamma > 0$. The focus was on the long-run steady-state in which continual series of entries and exits were observed. Consistent with the available empirical findings (Dunne *et al.*, 1988), the model generated persistent series of entries and exits. Also consistent with the empirical observations, the contemporary rates of entry and exit were shown to be positively correlated. Comparing the turnover rates across industries with different market-specific characteristics, I found that the mean rates of entry and exit move together across all industries: An industry with a higher-than-average rate of entry is also likely to have a higher-than-average rate of exit. Further delving into how market-specific factors affect the turnover rates of the firms and, consequently, the steady-state market structure, I found that the rate of firm turnover and the industry concentration are positively related. All of these results continue to hold with the present model where R&D is endogenous.

Both papers provided the foundation on which to build the base model of industry dynamics. The validation of the base model was made by generating results that were consistent with the empirical observations. Nevertheless, these earlier models were inadequate as the model of Schumpeterian competition, because the R&D process was specified as being exogenous. The main objective of the current paper is to endogenize the R&D decisions within the base model so as to properly study the relationship between R&D and firm turnovers in a unifying model of *creative destruction* as envisioned by Schumpeter.

3.4. Endogenizing the Process of R&D

> If we knew what it was we were doing, it would not be called research, would it? – Albert Einstein

The main goal of this paper is to endogenize the process of R&D into the base model of industry dynamics. This process corresponds to the stage-2 process of transforming $\underline{z}_i^{t-1}$ to $\underline{z}_i^t$ as described in Section 3.2.2. I model the R&D-related decisions as being driven by a set of choice probabilities that evolve over time on the basis of a reinforcement learning mechanism. If a firm decides to pursue R&D, it can do so through either *innovation* or *imitation*. The size of R&D expenditure depends on which of the two modes a given firm chooses: *Innovation* costs a fixed amount of K_{IN} while *imitation* costs K_{IM}. Hence, the necessary condition for a firm to engage in R&D is:

$$w_i^{t-1} \geq \max\{K_{IN}, K_{IM}\}.^{12} \tag{3.15}$$

Figure 2 illustrates the various stages of the R&D process. The crucial part of this model is how

[12] The computational experiments reported in this paper assume $K_{IN} > K_{IM}$.

the various components of the R&D decision are carried out. First, each firm i has two probabilities, α_i^t and β_i^t, which evolve over time via a reinforcement learning mechanism. Each period, firm i chooses to pursue R&D with probability α_i^t and not to pursue R&D with probability $1-\alpha_i^t$. If she chooses not to purse R&D, she simply keeps the old technology and, hence, $\underline{z}_i^t = \underline{z}_i^{t-1}$. However, if she chooses to pursue R&D, then she has a probability β_i^t with which she chooses to "innovate" and $1-\beta_i^t$ with which she chooses to "imitate." (As mentioned, both α_i^t and β_i^t are endogenous – how they are updated from one period to the next is discussed below.)

Innovation occurs when the firm considers changing the method (i.e., flipping the bit) in *one* randomly chosen activity. *Imitation* occurs when the firm (i) 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 (i's) current methods in all other activities.[13] 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:

$$p_{ij}^t \equiv \frac{\pi_j^{t-1}}{\sum_{\forall k \in S_*^{t-1}, k \neq i} \pi_k^{t-1}} \tag{3.16}$$

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 imitated.

Let $\widetilde{\underline{z}}_k^t$ 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:

$$\underline{z}_k^t = \begin{cases} \widetilde{\underline{z}}_k^t, & \text{if and only if } c_k(\widetilde{\underline{z}}_k^t, \widehat{\underline{z}}^t) < c_k(\underline{z}_k^{t-1}, \widehat{\underline{z}}^t) \\ \underline{z}_k^{t-1}, & \text{otherwise.} \end{cases} \tag{3.17}$$

Hence, a proposed technology is adopted by a firm if and only if it lowers the marginal cost below

[13] Hence, the imitating firm is capable of copying only a small part of the entire technology. This is one aspect of the cognitive limitation assumed in this research. An issue that can be investigated in the future is to relax this assumption and examine the impact that a firm's cognitive capacity has on the various outcomes at the firm and industry level. This is not pursued in this paper.

the level attained with the current technology the firm carries over from the previous period.[14] This happens when the Hamming distance to the optimal technology is lower with the proposed technology than with the current technology. Notice that this condition is equivalent to a condition on the firm profitability. When an incumbent firm takes all other incumbent firms' marginal costs as given, the only way that its profit is going to improve is if its marginal cost is reduced as the result of its innovation.

Note that firm i's R&D expenditure in period t depends on the type of R&D activity it pursued:

$$I_i^t = \begin{cases} 0 & \text{if no R\&D was pursued;} \\ K_{IN} & \text{if R\&D was pursued and innovation was chosen;} \\ K_{IM} & \text{if R\&D was pursued and imitation was chosen.} \end{cases} \tag{3.18}$$

Let us get back to the choice probabilities, α_i^t and β_i^t. Both probabilities are endogenous and specific to each firm. Specifically, they 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 course of action. The learning rule specifies how attractions are updated by the firm's experience and how the probabilities of choosing different courses of action depend on these attractions. The main feature 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, α_i^t and β_i^t are adjusted at the end of each period on the basis of evolving attraction measures: A_i^t for *R&D* and $\overline{A}_i^t$ for *No R&D*; B_i^t for *Innovation* and $\overline{B}_i^t$ for *Imitation*. Table 2 shows the adjustment dynamics of these attractions for the entire set of possible cases. According to this rule, A_i^t is raised by a unit when R&D (either through innovation or imitation) was productive and the generated idea was adopted. Alternatively, $\overline{A}_i^t$ is raised by a unit when R&D was unproductive and the generated idea was discarded. In terms of the choice between innovation and imitation, B_i^t is raised by a unit if R&D via innovation was performed and the generated idea was adopted or if R&D via imitation was performed and the generated idea was

[14] I assume that the evaluation of the technology by a firm in terms of its production efficiency (as representd by the level of its marginal cost) is done with perfect accuracy. While this assumption is clearly unrealistic, it is made to avoid overloading the model which is already substantially complicated.

discarded. Hence, the attraction for innovation can increase if either innovation was productive or imitation was unproductive. Conversely, $\overline{B}_i^t$ is raised by a unit if R&D via imitation generated an idea which was adopted – i.e., imitation was productive – or R&D via innovation generated an idea which was discarded – i.e., innovation was unproductive. If no R&D was performed, all attractions remain unchanged.

Given A_i^{t+1} and $\overline{A}_i^{t+1}$, one derives the choice probability of *R&D* in period $t+1$ as:

$$\alpha_i^{t+1} = \frac{A_i^{t+1}}{A_i^{t+1} + \overline{A}_i^{t+1}}. \tag{3.19}$$

In $t+1$, the firm then pursues *R&D* with probability α_i^{t+1} and *No R&D* with probability $1-\alpha_i^{t+1}$.

Given B_i^{t+1} and $\overline{B}_i^{t+1}$, one derives the choice probability of innovation in period $t+1$ as:

$$\beta_i^{t+1} = \frac{B_i^{t+1}}{B_i^{t+1} + \overline{B}_i^{t+1}}. \tag{3.20}$$

The probability of pursuing imitation is $1 - \beta_i^{t+1}$.

Finally, all new entrants in E^t are endowed with the initial attractions that make them indifferent to the available options. Specifically, I assume that $A_i^t = \overline{A}_i^t = 10$ and $B_i^t = \overline{B}_i^t = 10$ for a new entrant such that $\alpha_i^t = \beta_i^t = 0.5$ for all i – i.e., it has equal probabilities of choosing between *R&D* and *No R&D* as well as between *innovation* and *imitation*. Of course, these attractions will diverge from one another as the firms go through differential market experiences as the result of their R&D decisions made over time.

4. Design of Computational Experiments

The values of the parameters used in this paper, including those for the baseline simulation, are provided in Table 3.[15] I assume that there are 96 separate tasks in the production process, where the method chosen for each task is represented by a single bit. This implies that there are $2^{96} (\cong 8 \times 10^{28})$ different combinations of methods for the complete production process. In each period, there are exactly 40 potential entrants who consider entering the industry, where a new firm enters with a

[15] The source code for the computational experiments was written in C^{++} and the simulation outputs were analyzed and visualized using Mathematica 7.0. The source code is available upon request from the author.

start-up wealth (b) of 0. An incumbent firm will exit the industry, if his net wealth falls below the threshold rate ($\underline{W}$) of 0. The demand intercept is fixed at 300. Likewise, the cost of innovation, K_{IN}, is fixed at 100, while the cost of imitation, K_{IM}, is fixed at 50. All initial attractions for R&D activities are such that the new entrants are indifferent between *R&D* and *No R&D* as well as between *Innovation* and *Imitation*. The maximum magnitude of a change in technological environment, g, is 8 – that is, the Hamming distance between the optimal technologies at $t-1$ and at t can not be more than 8 bits. The time horizon is over 5,000 periods, where in period 1 the market starts out empty.[16]

The focus of my analysis is on the impacts of the market size (s) and the fixed cost (f) on the industry dynamics. I consider four different values for both parameters: $s \in \{4, 6, 8, 10\}$ and $f \in \{100, 200, 300, 400\}$. The impacts of s and f are also examined for four different rates of change in technological environment: $\gamma \in \{.1, .2, .3, .4\}$. Note that a higher value of γ reflects more frequent changes in technological environment. All other things held constant, this makes it tougher for firms to adapt to the changing environment.

Starting from an empty industry with the above configuration of parameters, I evolve the industry and trace its development by keeping track of the following endogenous variables:

- $|E^t|$: number of firms actually entering the industry in the beginning of t
- $|M^t|$: number of firms that are in operation in t (including both active and inactive firms)
- $|L^t|$: number of firms leaving the industry at the end of t
- $|S^t|$: number of firms surviving at the end of t ($= |M^t| - |L^t|$)
- P^t: market price at which goods are traded in t
- $\{c_i^t\}_{\forall i \in M^t}$: realized marginal costs of all firms that were in operation in t

[16] The examination of the simulation outputs shows that the horizon of 5,000 periods is more than enough for an industry to achieve a steady-state for all parameter values considered in this research.

- $\{q_i^t\}_{\forall i \in M^t}$: actual outputs of all firms that were in operation in t
- $\{\pi_i^t\}_{\forall i \in M^t}$: realized profits (losses) of all firms that were in operation in t
- $\{age_i^t\}_{\forall i \in M^t}$: ages of all firms in operation in t
- $\{\alpha_i^t\}_{\forall i \in M^t}$: *R&D* intensities of all firms that were in operation in t
- $\{\beta_i^t\}_{\forall i \in M^t}$: *innovation* intensities of all firms that were in operation in t

Using the above variables, I construct an additional group of endogenous variables that characterize the aggregate behavior of the firms in an industry. First, note that both the size of the market (s) and the fixed cost (f) are likely to have significant influence on the number of firms a given industry can sustain in the long run. Since the magnitude of firm turnovers must be viewed in relation to the size of the industry, I construct the *rates* of entry and exit, ER^t and XR^t, which are, respectively, the number of new entrants and the number of exiting firms as the fractions of the total operating firms in period t:

$$ER^t = \frac{|E^t|}{|M^t|} \text{ and } XR^t = \frac{|L^t|}{|M^t|}. \tag{4.1}$$

The rate of firm survival in period t is then $1 - XR^t$.

As a concentration measure, I use the Herfindahl-Hirschmann Index, H^t:

$$H^t = \sum_{\forall i \in M^t} \left(\frac{q_i^t}{\sum_{\forall j \in M^t} q_j^t} \cdot 100 \right)^2 . \tag{4.2}$$

To examine the intensity of the R&D activities, I look at the R&D expenditure per firm, ARD^t:

$$ARD^t = \frac{\sum_{\forall i \in M^t} I_i^t}{|M^t|}.$$

In any given period, the aggregate R&D expenditure, $\sum_{\forall i \in M^t} I_i^t$, consists of the amount spent by the firms that innovate and the amount spent by those that imitate. Let NTM^t be the ratio of the two components:

$$NTM^t = \frac{\text{total innovation expenditure in } t}{\text{total imitation expenditure in } t} \tag{4.3}$$

It, hence, measures the firms' tendency to innovate rather than to imitate each other.

For an aggregate measure of the industry's production efficiency, I construct an industry marginal cost, WMC^t, where

$$WMC^t = \sum_{\forall i \in M^t} \left[\left(\frac{q_i^t}{\sum_{\forall j \in M^t} q_j^t} \right) \cdot c_i^t \right]. \tag{4.4}$$

WMC^t is, hence, the weighted average of the individual firms' marginal costs in period t, where the weights are the market shares of the firms in that period.

Finally, I construct an aggregate measure of firms' price-cost margins, PCM^t, where

$$PCM^t = \sum_{\forall i \in M^t} \left[\left(\frac{q_i^t}{\sum_{\forall j \in M^t} q_j^t} \right) \cdot \left(\frac{P^t - c_i^t}{P^t} \right) \right]. \tag{4.5}$$

PCM^t is the weighted average of the individual firms' price-cost margins in period t, where the weights are the market shares of the operating firms in that period.

In Section 5.1, I examine the time series values of a subset of these endogenous variables from a single replication based on the baseline parameter values.[17] The properties reported in the section are significant in that they were observed for many independent replications, each using a fresh set of random numbers. The comparative study carried out in Section 5.2 is based on 500 independent replications for each parameter configuration. The steady-state values (average over the periods from 3001 to 5000) of the endogenous variables at the industry-level were then averaged over those 500 replications. The resulting mean steady-state values for the relevant endogenous variables are denoted as follows: $\overline{ER}$, $\overline{XR}$, $\overline{H}$, $\overline{P}$, $\overline{ARD}$, $\overline{NTM}$, $\overline{WMC}, \overline{PCM}$. These values are compared for different parameter configurations representing different industries.

5. Results

5.1. Patterns across Time and across Firms

The first step in my analysis is to examine the heterogeneity among firms that emerges over time, as a typical industry develops and matures through a continual series of entries and exits by firms.

[17] A single replication is a simulation run of the previously described 4-stage process repeated over the first 5,000 periods.

A "typical" industry is specified by the baseline parameter values as indicated in Table 3. I start by observing the endogenous time paths of the three turnover variables over the 5,000 periods of the industry's development from its birth to full maturity: (a) the number of entrants, $|E^t|$; (b) the number of exiting firms, $|L^t|$; and (c) the number of operating firms, $|M^t|$. The results are captured in Figure 3 – these are from a single replication using the baseline parameter values.

Note from the top of Figure 3 the initial surge in the number of new entrants into the industry at its birth: The entire pool of potential entrants (40) jumps into the industry as it is newly born. This rush quickly slows down and the industry settles into a steady state where there are occasional entries that continue indefinitely over the horizon. The middle figure shows that the initial surge of entries is immediately followed by a large number of exits, implying that a large number of firms who initially entered the industry are soon forced out through a severe market competition – i.e., a "shakeout.". After the initial shakeout, the industry experiences a steady out-flow of firms that accompanies the steady in-flow of firms exhibited in the top figure. Hence, the industry experiences a *persistent* series of entry and exit. The continual streams of entries and exits interact to produce the time series in the bottom figure on the total number of operating firms, $|M^t|$, which include both the active and inactive firms. The time path shows that the number of operating firms moves with substantial volatility over time, though it moves around a steady mean ($\approxeq 43$) after about t = 1000. This is a natural consequence of $|E^t|$ and $|L^t|$ being positively correlated for all t. For the baseline run reported in Figure 3, the correlation was 0.6.[18] The positive correlation between the numbers of entries and the numbers of exits holds for all other runs tried in this study.[19]

The time paths captured in Figure 3 are typical of all replications performed in this study. For multiple replications using the same baseline values for parameters but with fresh random numbers, it is shown that the distribution of the outcomes generated from the stochastic process tends to be time-invariant for $t > 1,000$. Figure 4 shows the time series outputs of the same variables as in

[18] The correlation between the *rates* of entry and exit (defined as the number of entry or exit over the total number of operating firms) was 0.58 for the baseline run.

[19] This property is more fully described in Chang (2011), which used the same base model, but with exogenous and costless R&D.

Figure 3 as the mean over 100 independent replications: It is clear that the number of operating firms on average attains a stable level by $t = 1,000$. In fact, the time paths of interest always reach a steady state by $t = 3,000$ for all parameter configurations considered in this study.[20] As such, when we examine the impact of industry-specific factors on the industry's performance in Section 5.2, the steady-state value of an endogenous variable will be computed as an average over the last two thousand periods between $t = 3,001$ and $t = 5,000$.

The persistence of firm entries and exits over time comes from the unexpected shifts in the technological environment surrounding the firms (which happens at the rate of γ). To see this, I track the occurrences of technological shifts over the entire horizon. For a given technological shift that occurs in period τ, I define its "episode" as those consecutive periods following the shift before the next technological shift occurs in period τ'. The duration of the episode is then $(\tau' - \tau)$. For the baseline run captured in Figure 3, there was a total of 502 episodes of varying durations. Figure 5(a) shows the frequencies of different episode durations for the technological shifts that occurred from $t = 1,001$ to $5,000$.[21]

To see the impact that technological shifts have on the turnover of firms, I ask, for each period over the horizon, how many periods have elapsed since the last technological shift. This allows me to examine the relationship between the rates of entry and exit and the elapsed time since a technological shift. Figures 5(b) and 5(c) capture this information (again from $t = 1,001$ to $5,000$). It is clear that both rates tend to fall as the given period is further away from the last technological shift. In the presence of continual technological shifts, one may then view the turnover dynamics as being a series of *mini-shakeouts*, in which the rates of entry and exit jump up immediately following a technological shift and then gradually fall down as the market adjusts to the new environment. The positive correlation between the rates of entry and exit is then a natural consequence of this repeated shakeouts following technological shifts.

Each operating firm has three endogenous variables that are particularly relevant for our analy-

[20] See Law and Kelton (2000) for detailed discussions on how to identify the steady state in stochastic processes.
[21] The initial $1,000$ periods are ignored, given the transient nature of the infant industry.

sis: the firm's age (age_i^t), its marginal cost (c_i^t), and the intensity with which it pursues R&D (α_i^t). For any given industry, we can explore the evolving relationships between: 1) the distribution of firms' ages within the industry; 2) the distribution of marginal costs within the industry; and 3) the distribution of R&D intensities within the industry.

Figure 6 presents the snapshots of the relationship between the ages of the firms and their R&D intensities at twelve different points in time over the first 5,000 periods. These are from the single replication using baseline parameter values – i.e., the same run that generated Figure 3. At $t = 1$, all new entrants enter with $\alpha_i^0 = 0.5$ such that they are indifferent between pursuing and not pursuing R&D. Once in the industry, these probabilities tend to adjust downward but at different rates based on differences in the firms' experiences. By the time they reach $t = 100$, we already observe a strong negative correlation between age_i^t and α_i^t, where an older firm tends to pursue R&D with a lower intensity than a more recent arrival. The intuition is that with occasional technological shifts there are persistent entries by new firms who tend to enter with higher propensity for R&D than the older incumbents. Their coexistence in the market then leads to a negative correlation between the ages and the R&D intensities of the firms. The following property emerges and remains visible once the industry enters the steady state.

Property 1: Within an industry, the R&D intensity of a firm is *negatively* related to its age.

Given the general weakening of the R&D intensity over a firm's life, a natural question is what happens to its level of production efficiency as measured by the marginal cost, c_i^t. Do firms become relatively inefficient as they age? This evolving relationship between the age of a firm and its marginal cost is captured in Figure 7 and summarized in Property 2.

Property 2: Within an industry, the marginal cost of a firm is *positively* related to its age.

Recall that the new entrants enter with heterogeneous technologies randomly selected from the technology space. The initial distribution of marginal costs is likely to be widely variant. Over time, there will be two countervailing forces acting on this distribution: 1) the selection effect of market competition, which tends to reduce the variability by weeding out the inefficient firms; 2) the turbulence in the technological environment, which increases the diversity of technologies by creating opportunities for new entrants to come in with technologies better suited for the new environment. Figure 7 shows that over time we tend to observe a positive relationship between the firm age and the level of its marginal cost: Even though the selection effect eventually weeds out the old firms that are unable to improve their production efficiency over time, the initial accumulation of their wealth at the early stages of their lives allow them to linger on for a sufficiently long period of time so that we observe efficient young firms and inefficient old firms coexisting side by side.

Properties 1 and 2 jointly imply a *negative* relationship between the marginal cost of a firm and its R&D intensity: Those firms that are more active in R&D are more efficient in production. Since a firm's operating efficiency is captured in the form of its marginal cost of production, it also follows that younger firms with relatively lower marginal costs tend to have bigger market shares and higher sales revenues. This predicts a positive relationship between the R&D intensity of a firm and its size (as represented by the firm's sales revenue, $P^t \cdot q_i^t$). The outcome generated by the computational model confirms this intuition – see Figure 8.

Property 3: Within an industry, the R&D intensity of a firm and its sales revenue are *positively* related.

Property 3 then confirms the first of the two Schumpeterian hypotheses that bigger firms tend to be more active in R&D. I am, however, unable to identify any significant correlation between these endogenous variables and β_i^t, the propensity to pursue innovation vs. imitation.

Although Figures 6-8 only show a few snapshots of the relationships reported in Properties 1-3, the complete time series of the correlations between the firm-level endogenous variables confirm that they hold consistently for all periods over the relevant horizon.[22] Also, note that the simulation

[22] The average correlations over the last 2,000 periods from $t = 3,001$ to $5,000$ were: 1) $-.664589$ between the age and the R&D intensity; 2) .528621 between the age and the marginal cost; and 3) .536242 between the R&D intensity and the sales revenue.

results that give rise to the above properties are based on the baseline values of $\alpha_i^0 = \beta_i^0 = 0.5$. In the absence of any information about the propensities of new entrants to perform R&D, this specification seems reasonable. Alternatively, one could consider α_i^0 and β_i^0 to be randomly chosen from [0, 1] for each i based on uniform distribution. The simulations with this specification (not reported here) generated results that are qualitatively identical to those reported in this paper, including the above three properties.

5.2. Patterns across Industries

Note that there are three important parameters (industry-specific factors) in this model; the market size (s), the fixed cost (f), and the rate of change in technological environment (γ). The main objective in this section is to examine how these parameters affect the long-run development of an industry. As mentioned earlier, I focus on the steady-state values of the endogenous variables averaged over 500 independent replications for each parameter configuration that represents a particular industry.

5.2.1. Comparative Dynamics Results

Let us start our analysis by first looking at the above endogenous variables for various combinations of the structural parameters, s and f, given a fixed value of γ. Fixing γ at the baseline value of .1, I plot in Figures 9 and 10 the rate of entry ($\overline{ER}$) and the rate of exit ($\overline{XR}$) for all $s \in \{4, 6, 8, 10\}$ and $f \in \{100, 200, 300, 400\}$. They show that both rates decline with the size of the market, s, but increase with the fixed cost, f. Since the two rates move in the same direction in response to the changes in the two parameters, I will say that the rate of *firm turnover* increases or decreases as either of these rates goes up or down, respectively.[23]

What are the implications of these entry-exit dynamics on the structure of the industry and how are they affected by the market size and the fixed cost? Figure 11 shows that the steady-state mean number of operating firms increases with the size of the market and decreases with the size

[23] Some authors in the past have used the sum of the two rates as the rate of firm turnover. Given that the two rates move in the same direction, all of the results in this paper are consistent with that definition.

of the fixed cost. Consistent with this result, Figure 12 shows that the industry concentration, as measured by HHI ($\overline{H}$), decreases with the size of the market and increases with the size of the fixed cost.

We also find that the steady-state market price of an industry is *negatively* related to the size of the market and *positively* related to the fixed cost – see Figure 13. The properties on industry concentration and market price are clearly due to the degree of market competition as endogenously determined by the market size and the fixed cost – that is, the degree of competition is higher in larger markets and/or in markets with lower fixed costs.

Moving on to the R&D activities, Figures 14 looks at $\overline{ARD}$ which shows how much a single firm spends on R&D on average. The following property is observed: The steady-state average R&D expenditure per firm is *negatively* related to the size of the market and *positively* related to the fixed cost. An implication of this result is the level of marginal costs attained by the firms as a function of these parameters. Figure 15 captures the steady-state mean of the industry marginal cost, $\overline{WMC}$: The industry marginal cost is *negatively* related to market size and *positively* related to the fixed cost.

Further delving into the R&D activities of the firms, I now compute, $\overline{NTM}$, the ratio of the aggregate innovation expenditure to the aggregate imitation expenditure and track its evolution over time. Figure 16 shows the steady-state mean of this ratio is *negatively* related to the size of the market and *positively* related to fixed cost. In other words, firms operating in a larger market or having lower fixed costs are more likely to invest in imitative R&D than innovative R&D and vice versa. Combined with the result on market concentration (as reported in Figure 12), this also implies that the firms in more (less) concentrated markets are more (less) likely to pursue innovative R&D than imitative R&D.

Finally, Figure 17 shows the industry price-cost-margin ($\overline{PCM}$) is *negatively* related to market size and *positively* related to fixed cost. Note that both price ($\overline{P}$) and the industry marginal cost ($\overline{WMC}$) are similarly affected by the market size and fixed cost: A smaller market or a higher fixed cost leads to higher price and higher marginal cost. But it is the extra gain in price coming from the increased market power for firms that dominates the simultaneous increase in marginal costs.

The comparative dynamics results are then summarized as follows:

Property 4: The steady-state values of the following endogenous variables are *negatively* related to the size of the market and *positively* related to fixed cost: 1) the rate of firm turnovers; 2) industry concentration; 3) market price; 4) R&D spending per firm; 5) industry marginal cost; 6) ratio of innovation-to-imitation spending per firm; 7) industry price-cost-margin.

Note that the size of the market and the fixed cost influence the seven endogenous variables – $\overline{ER}$ ($\overline{XR}$), $\overline{H}$, $\overline{P}$, $\overline{ARD}$, $\overline{WMC}$, $\overline{NTM}$, and $\overline{PCM}$ – in a uniform way. It is straightforward to establish the relationships between the endogenous variables on the basis of Property 4. For instance, an industry with a high rate of firm turnover is likely to be highly concentrated. The firms in such an industry are likely to invest heavily in R&D, with a greater emphasis on innovation than imitation. The same industry tends to have higher industry marginal cost (and thus be relatively inefficient), but generate higher price-cost margins for the firms. These predicted relationships have implications for cross-sectional empirical research as discussed in the next section.

5.2.2. Implications for Cross-Industries Studies

Most of the empirical studies in industrial organization addressing the issues of market structure and performance are cross-sectional studies of a large number of heterogeneous industries. These industries are likely to vary widely in terms of the size of the market they face as well as the average fixed cost that determines the economies of scale for the firms within each industry. Given the diverse sample of industries, the past empirical studies attempted to identify relationships between variables that are endogenous to the industry dynamics, such as the rate of firm turnover, industry concentration, and R&D intensities. The comparative dynamics results presented in the previous section can be useful in understanding these cross-sectional results within a unifying conceptual framework.

Suppose the industries in the sample are differentiated in terms of s and f, where $s \in \{4, 6, 8, 10\}$ and $f \in \{100, 200, 300, 400\}$. Such a collection of industries will then display positive relationships between the seven observed endogenous variables: rate of firm turnover, concentration, per-firm R&D spending, industry marginal cost, innovation-to-imitation spending ratio, and industry price-cost margin.[24] Most importantly, the following property emerges from these results.

Property 5: The firms in a more (less) concentrated industry tend to spend more (less) on R&D per firm.

The positive relationship between the degree of concentration and the R&D intensity confirms the second Schumpeterian Hypothesis, if we accept the average R&D per firm as the valid measure of R&D intensity.

Other relevant results that can be inferred from Property 4 include: 1) firms in more concentrated industries are less efficient (i.e., have higher industry marginal cost), but they have higher price-cost margins; 2) firms in more concentrated industries tend to spend more on innovation than on imitation. The first result implies a positive relationship between concentration and price-cost margin, but the source of the high price-cost margin in this case is the market power (higher price) rather than efficiency. The second property – firms in concentrated industry are more innovative, while those in unconcentrated industry are more imitative – may be due to the fact that the return to imitation is larger in less concentrated industries, where there exists a larger and more diverse set of sources from which a firm can copy its practices from. A concentrated industry with relatively small number of firms does not provide such a large pool of proven ideas and a firm has to rely on innovation to a greater extent.

[24] In fact, one could directly simulate this scenario by creating an artificial population of industries, each having random draws for s and f from the pre-specified ranges of values. The model presented here can then be run for each industry separately. One can perform the standard regression analyses on the steady-state outcomes from these industries, mimicking the actual empirical studies. The results from this computational experiment should, of course, be consistent with those reported in this paper.

5.3. Robustness and the Impact of the Rate of Technological Change (γ)

Properties 1-5 were established given $\gamma = .1$. The same sets of simulations were carried out for $\gamma \in \{.2, .3, .4\}$ to check for robustness. All of the above properties continue to hold for these values of γ. In addition, I am able to identify additional properties on the impact of γ on various endogenous variables. Some of these properties are shown in Figure 18.[25] The rate of entry (and the rate of exit, not shown here) increase with γ. The degree of concentration (HHI) is independent of γ: while γ influences the rates of entry and exit, the positive and high correlation between the two rates implies that the degree of concentration is insensitive to γ. The market price increases with γ. The industry marginal cost also increases with γ as a technologically more turbulent environment makes it more difficult to find the efficient technology. The industry price-cost margin decreases with γ: while both the price and the marginal costs increase with γ, the rise in marginal cost dominates and, hence, the squeezing of the price-cost margin. In terms of the R&D efforts, the average R&D expense per firm increase with γ.

Note that the R&D expenditure per firm consists of innovation expenditure and imitation expenditure per firm. The bottom right figure in Figure 18 captures the ratio of innovation expenditure to R&D expenditure per firm for varying degrees of technological turbulence. It shows that the share of the innovation cost in a firm's R&D cost tends to rise with γ, hence implying that a more turbulent technological environment leads to more innovative R&D than imitative R&D. This result comes from the distinct ways in which innovation and imitation affect the evolving diversity of the technologies held by the firms. Note that a firm, once it chooses to pursue R&D in stage 2 of its decision making, faces two options for carrying this out, innovation and imitation. From an individual firm's perspective, the act of innovation entails trying out a random idea (technology) taken from the entire pool of ideas (the whole technology space). The act of imitation, on the other hand, entails trying out ideas from a restricted subset which largely contains those ideas that have

[25] The box in these plots spans the distance between two quantiles surrounding the median with lines that extend to span the full dataset.

been tried by other firms and found useful. To the extent that imitation considers proven ideas only, it would seem to be more effective than innovation from an individual firm's perspective. However, this seemingly obvious property must be considered in conjunction with the fact that the technological environment is volatile in the proposed model: What was useful in one environment may no longer be useful in the new environment – recall that the optimal technology, which is the target of R&D search, can change from one period to the next. Consequently, the advantage to a single firm of imitating relative to innovating depends on the degree of volatility in the industry's technological environment; with a more volatile environment favoring the innovation strategy and a more stable environment favoring the imitation strategy. The computational result presented in Figure 18 confirms this intuition.

6. Concluding Remarks

Viewed in the context of Schumpeter's "perennial gale of creative destruction," it is clear that a coherent understanding of the persistent firm turnovers and the associated regularities is possible only when they are considered in connection with the innovative efforts of firms to adapt and survive on the face of changing economic conditions. The model presented in this paper is constructed to incorporate the central features of the Schumpeterian process into a general model of industry dynamics. It generated persistent heterogeneity among firms through entry and exit induced by the unexpected shocks in the technological environment. The heterogeneity in firm ages coupled with the endogenous R&D carried out through experiential learning generated a positive relationship between firm size and the R&D intensity within industry. Allowing industries to differ in terms of the size of market demand and the fixed cost, I was further able to study the relationships between the relevant endogenous variables such as the rate of firm turnovers, industry concentration, and R&D expenditure per firm. The cross-industry comparisons generated a number of interesting results; one of the results was the positive relationship between the industry concentration and the R&D expenditure per firm, confirming the second Schumpeterian hypothesis.

While this work fits into the evolutionary and computational approach to industry dynamics, it is unique in several respects. First, it explicitly models the process of entry and exit while simultaneously endogenizing the process of firm R&D through a reinforcement-learning mechanism. This allows a coherent understanding of the two separately-developed literatures on R&D and firm turnover. Second, it is capable of generating industry dynamics with a realistically large number of firms, hence potentially providing a better fit to empirical data. Third, the computational approach utilized in this research allows for comparative dynamics studies, substantive enough to identify the industry-specific factors affecting the evolving structure and performance of the industries. These singular features result in a very rich conception of the innovative behavior of the firm and its technological and competitive environment that ultimately leads to substantive understanding as to the evolutionary dynamics of industries.

References

[1] Camerer, C. and T.-H. Ho, "Experience-Weighted Attraction Learning in Normal Form Games," *Econometrica*, 67 (1999), 827-874.

[2] Camerer, C. and D. Lovallo, "Overconfidence and Excess Entry: An Experimental Approach," *American Economic Review*, 89 (1999), 306-318.

[3] Caves, R. E., "In Praise of the Old I.O.," *International Journal of Industrial Organization*, 25 (2007), 1-12.

[4] Chang, M.-H., "Industry Dynamics with Knowledge-Based Competition: A Computational Study of Entry and Exit Patterns," *Journal of Economic Interaction and Coordination*, 4 (2009), 73-114.

[5] Chang, M.-H., "Entry, Exit, and the Endogenous Market Structure in Technologically Turbulent Industries," *Eastern Economic Journal*, 37 (2011), 51-84.

[6] Cohen, W. M. and R. C. Levin, "Empirical Studies of Innovation and Market Structure," in *Handbook of Industrial Organization, Volume II*, edited by R. Schmalensee and R. D. Willig, Elsevier Science Publishers B.V., 1989.

[7] Doraszelski, U. and A. Pakes, "A Framework for Applied Dynamic Analysis in IO," in *Handbook of Industrial Organization*, edited by M. Armstrong and R. Porter, Volume 3, Elsevier, Amsterdam, 2007.

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

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

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

[11] Law, A. M. and D. Kelton, *Simulation Modeling and Analysis*, 3rd ed., McGraw-Hill, New York, NY, 2000.

[12] Nelson, R. R. and S. G. Winter, *An Evolutionary Theory of Economic Change*, Harvard University Press, Cambridge, MA, 1982.

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

[14] Schumpeter, J. A., *Capitalism, Socialism, and Democracy*, 3rd ed., Harper, New York, NY, 1950.

[15] Tesfatsion, L. and K. Judd, eds., *Handbook of Computational Economics, Vol. 2: Agent-based Computational Economics*, Handbooks in Economics Series, North Holland, Amsterdam, 2006.

[16] Weintraub, G. Y., C. L. Benkard, and B. Van Roy, "Markov Perfect Industry Dynamics with Many Firms," *Econometrica*, 76 (2008), 1375-1411.

[17] Weintraub, G. Y., C. L. Benkard, and B. Van Roy, "Computational Methods for Oblivious Equilibrium," *Operations Research*, 58 (2010), 1247-1265.

Table 1 : Set Notations

Notation	Definition
S^t	Set of surviving firms at the end of t
S_*^t	Those in S^t which were profitable 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 ($= S^{t-1} \cup E^t$)
L^t	Set of firms which exit the industry at the end of t

Table 2 : Evolving Attractions

Decision Path			Updating of Attractions			
No R&D			$A_i^{t+1}= A_i^t$;	$\overline{A}_i^{t+1}=\overline{A}_i^t$;	$B_i^{t+1}= B_i^t$;	$\overline{B}_i^{t+1}=\overline{B}_i^t$
R&D	Innovation	Adopt	$A_i^{t+1}= A_i^t+1$;	$\overline{A}_i^{t+1}=\overline{A}_i^t$;	$B_i^{t+1}= B_i^t+1$;	$\overline{B}_i^{t+1}=\overline{B}_i^t$
		Discard	$A_i^{t+1}= A_i^t$;	$\overline{A}_i^{t+1}=\overline{A}_i^t+1$;	$B_i^{t+1}= B_i^t$;	$\overline{B}_i^{t+1}=\overline{B}_i^t+1$
	Imitation	Adopt	$A_i^{t+1}= A_i^t+1$;	$\overline{A}_i^{t+1}=\overline{A}_i^t$;	$B_i^{t+1}= B_i^t$;	$\overline{B}_i^{t+1}=\overline{B}_i^t+1$
		Discard	$A_i^{t+1}= A_i^t$;	$\overline{A}_i^{t+1}=\overline{A}_i^t+1$;	$B_i^{t+1}= B_i^t+1$;	$\overline{B}_i^{t+1}=\overline{B}_i^t$

Table 3 : List of Parameters and Their Values

Notation	Definition	Baseline Value	Parameter Values Considered
N	Number of tasks	96	96
r	Number of potential entrants per period	40	40
b	Start-up wealth for a new entrant	0	0
$\underline{W}$	Threshold level of net wealth for survival	0	0
a	Demand intercept	300	300
f	Fixed production cost	200	{100, 200, 300, 400}
K_{IN}	Fixed cost of innovation	100	100
K_{IM}	Fixed cost of imitation	50	50
A_i^0	Initial attraction for R&D (all i)	10	10
$\overline{A}_i^0$	Initial attraction for No R&D (all i)	10	10
B_i^0	Initial attraction for Innovation (all i)	10	10
$\overline{B}_i^0$	Initial attraction for Imitation (all i)	10	10
g	Maximum magnitude of change in technological environment	8	8
T	Time horizon	5,000	5,000
s	Market size	4	{4, 6, 8, 10}
γ	Rate of change in technological environment	0.1	{0.1, 0.2, 0.3, 0.4}

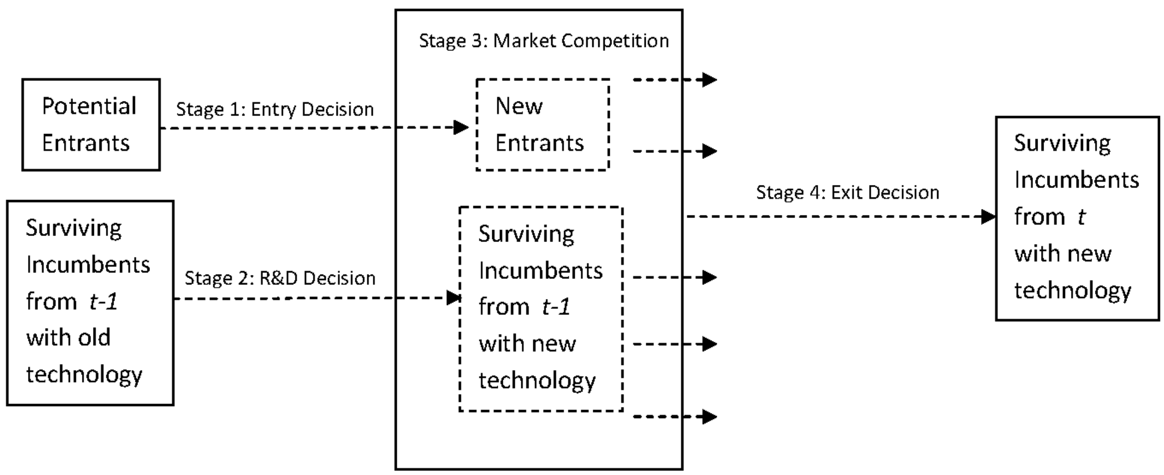

Figure 1: Four Stages of Decision Making in Period *t*

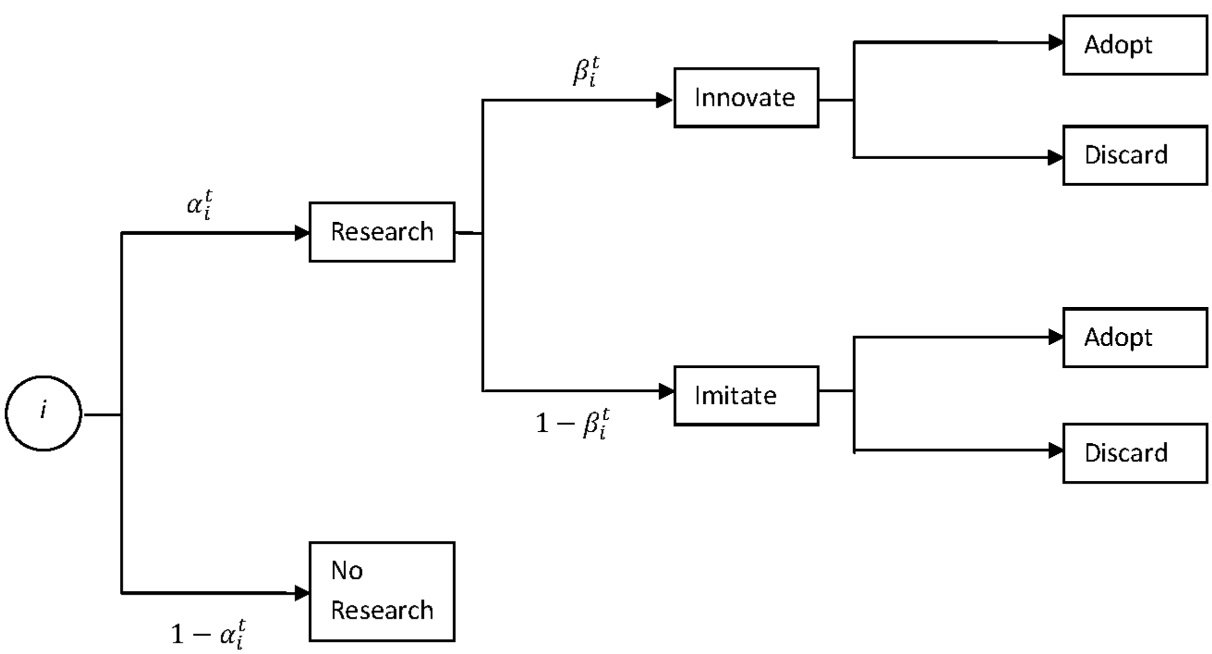

Figure 2: R&D Decision in Stage 2

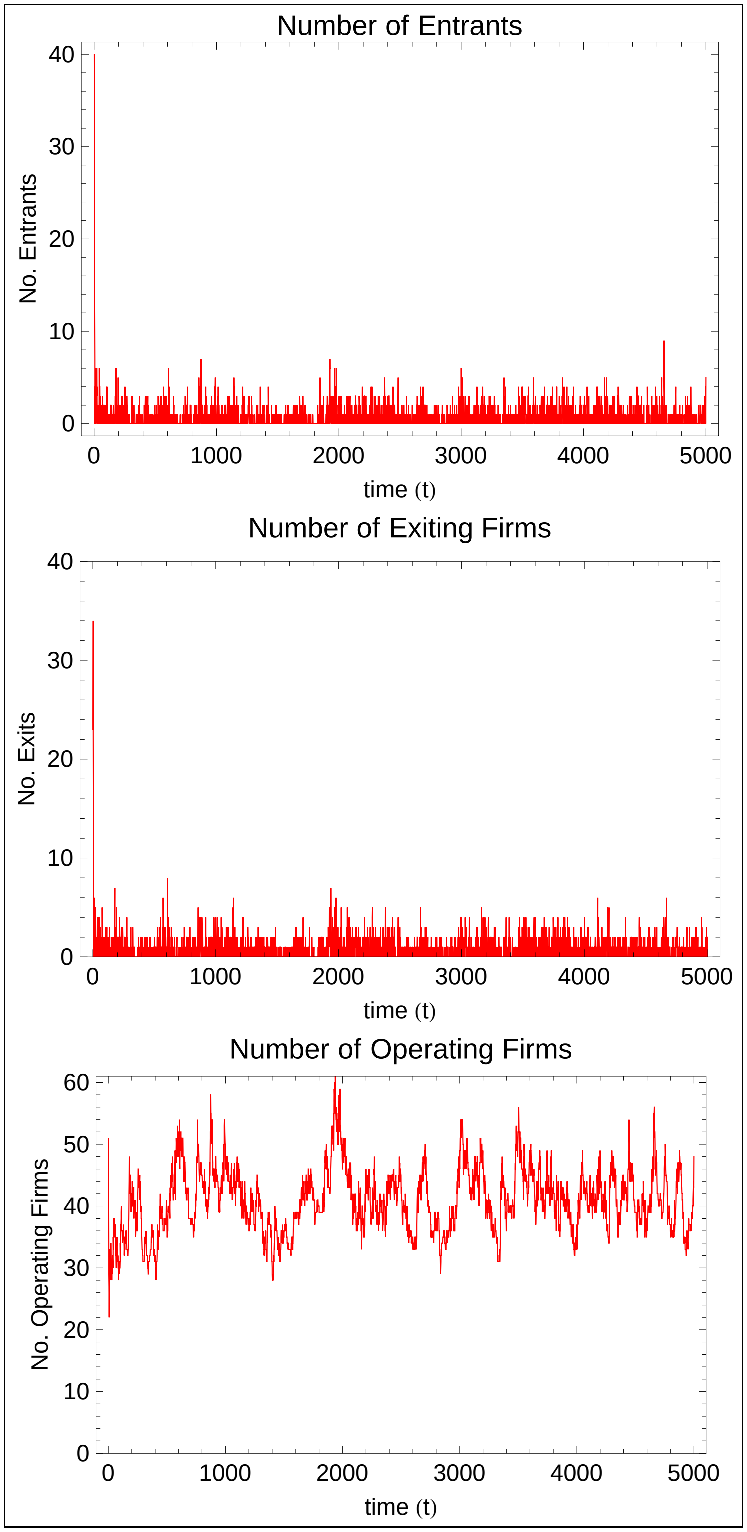

Figure 3: Endogenous Turnovers of Firms

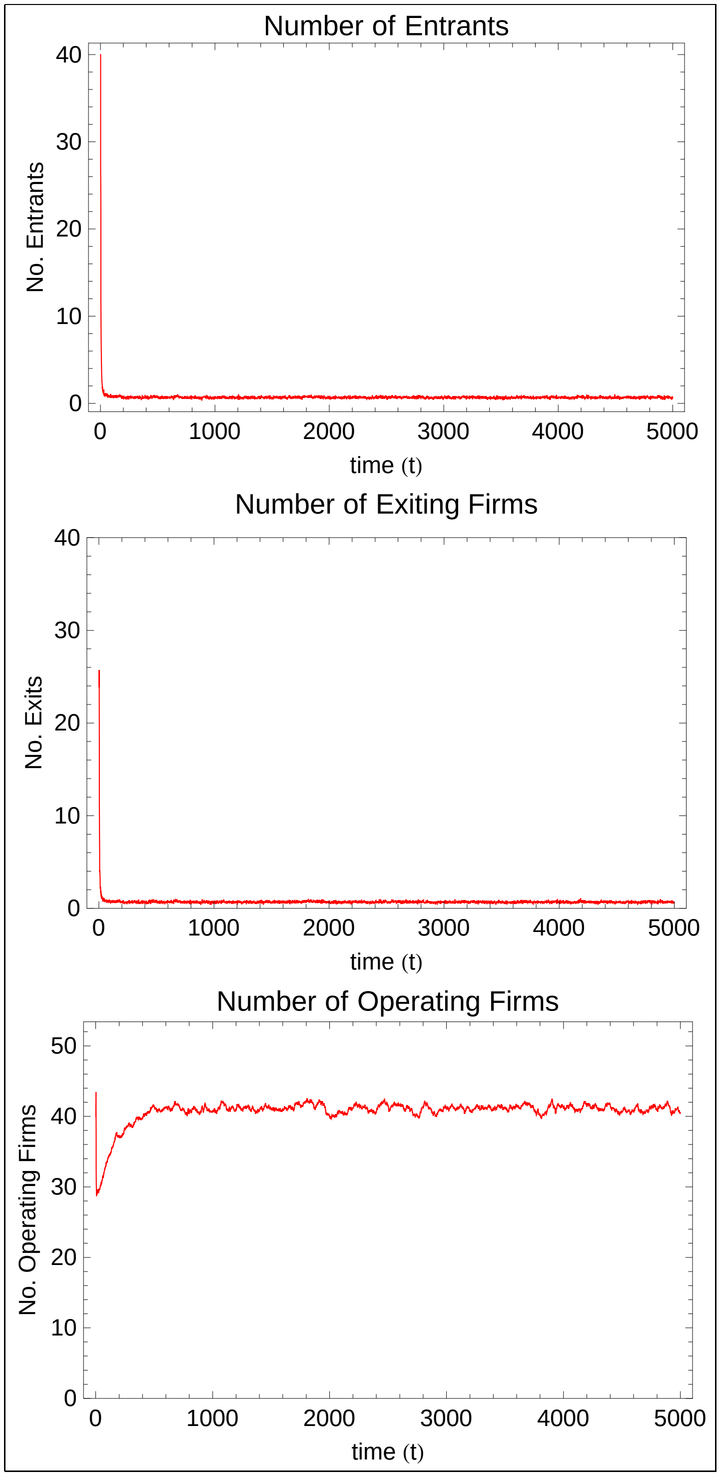

Figure 4: Endogenous Turnovers of Firms
(average over 100 replications)

(a)

(b)

(c)

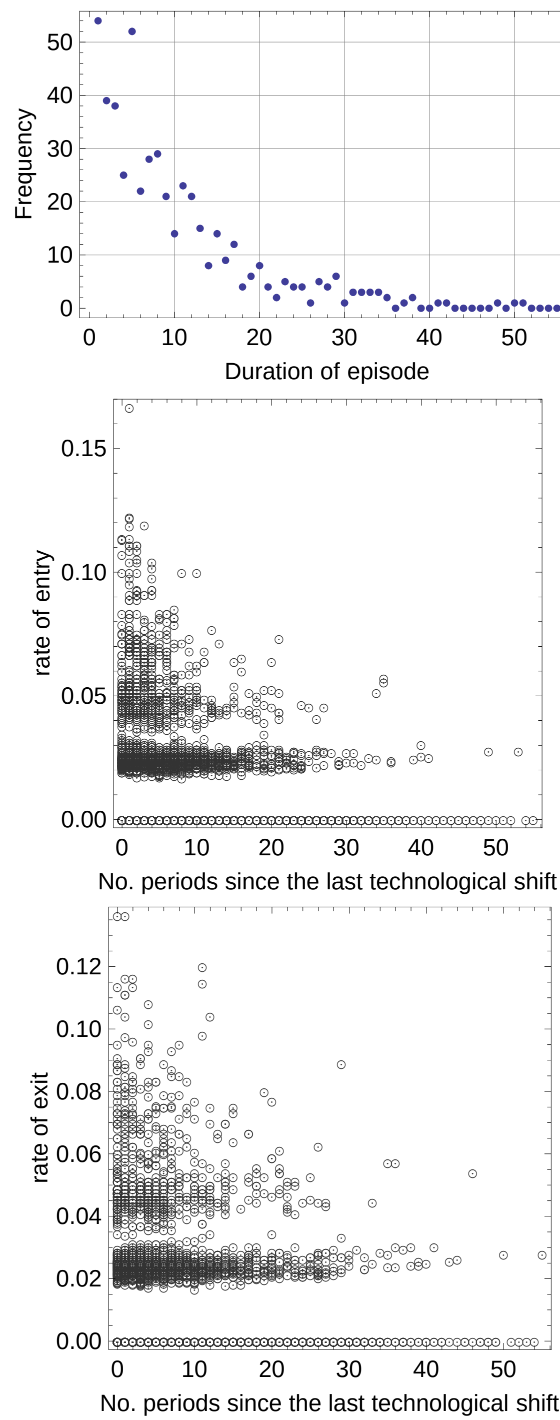

Figure 5: Technological Shifts and Turnovers

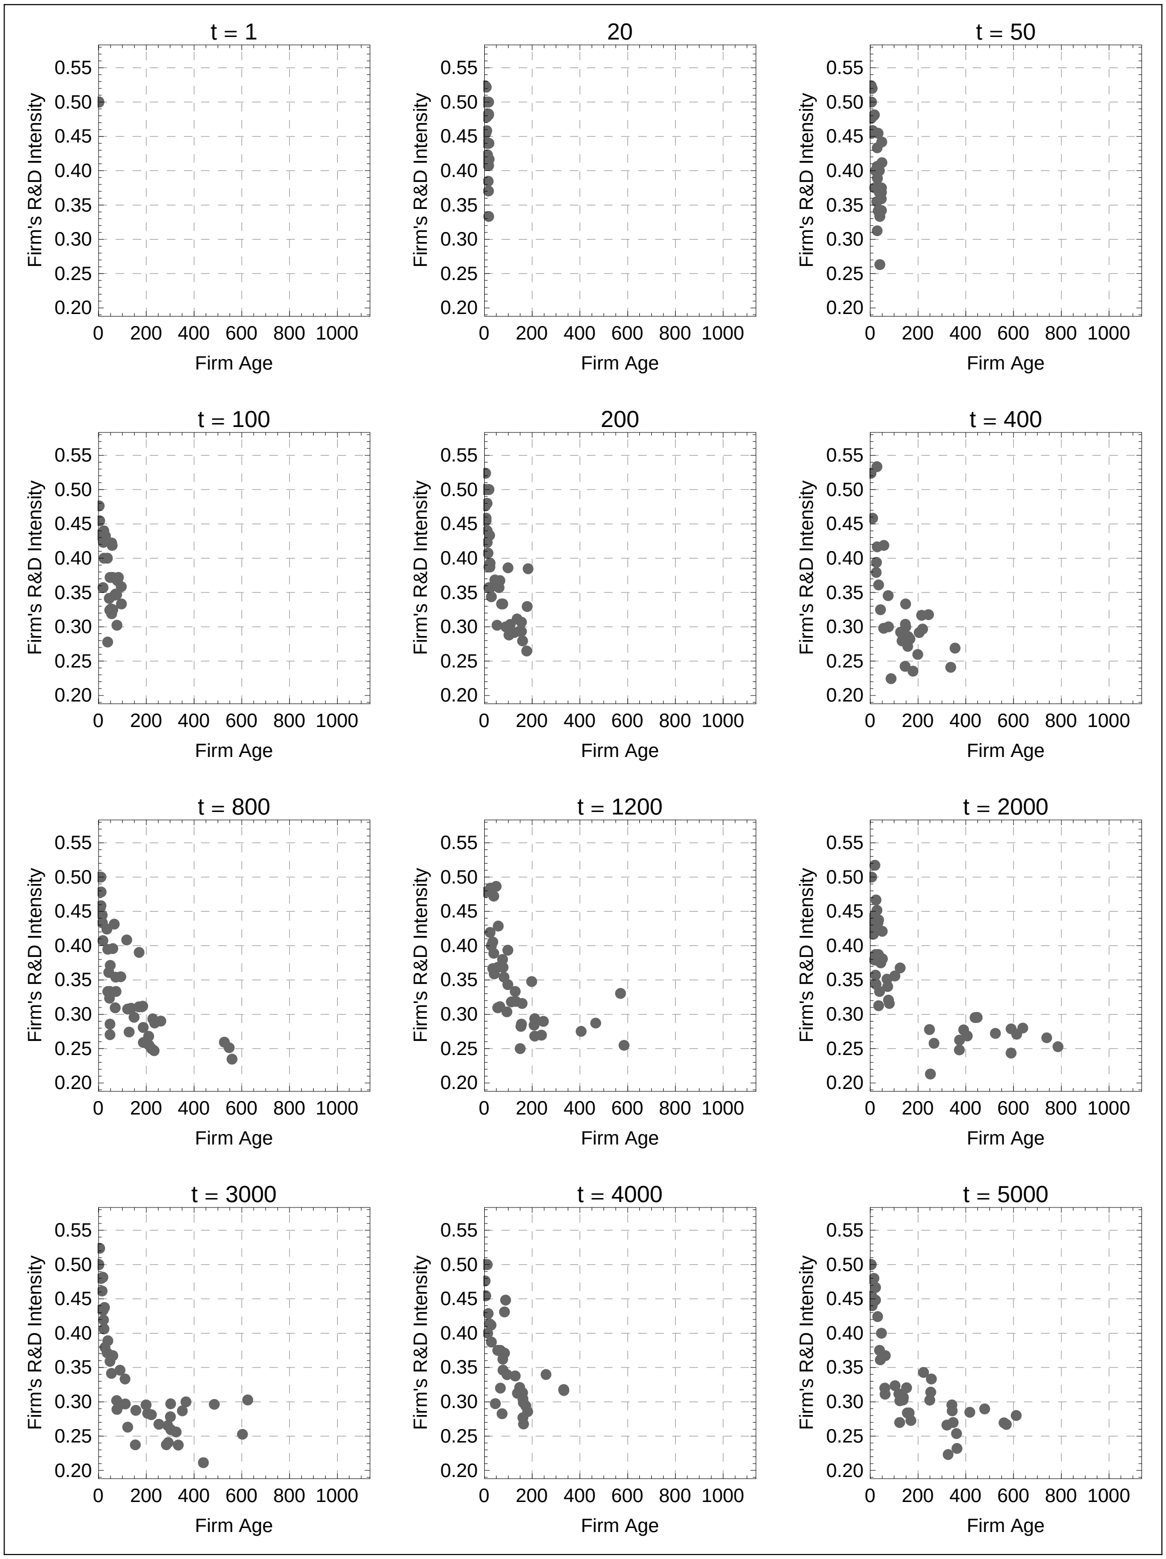

Figure 6: Firm Age and R&D Intensity

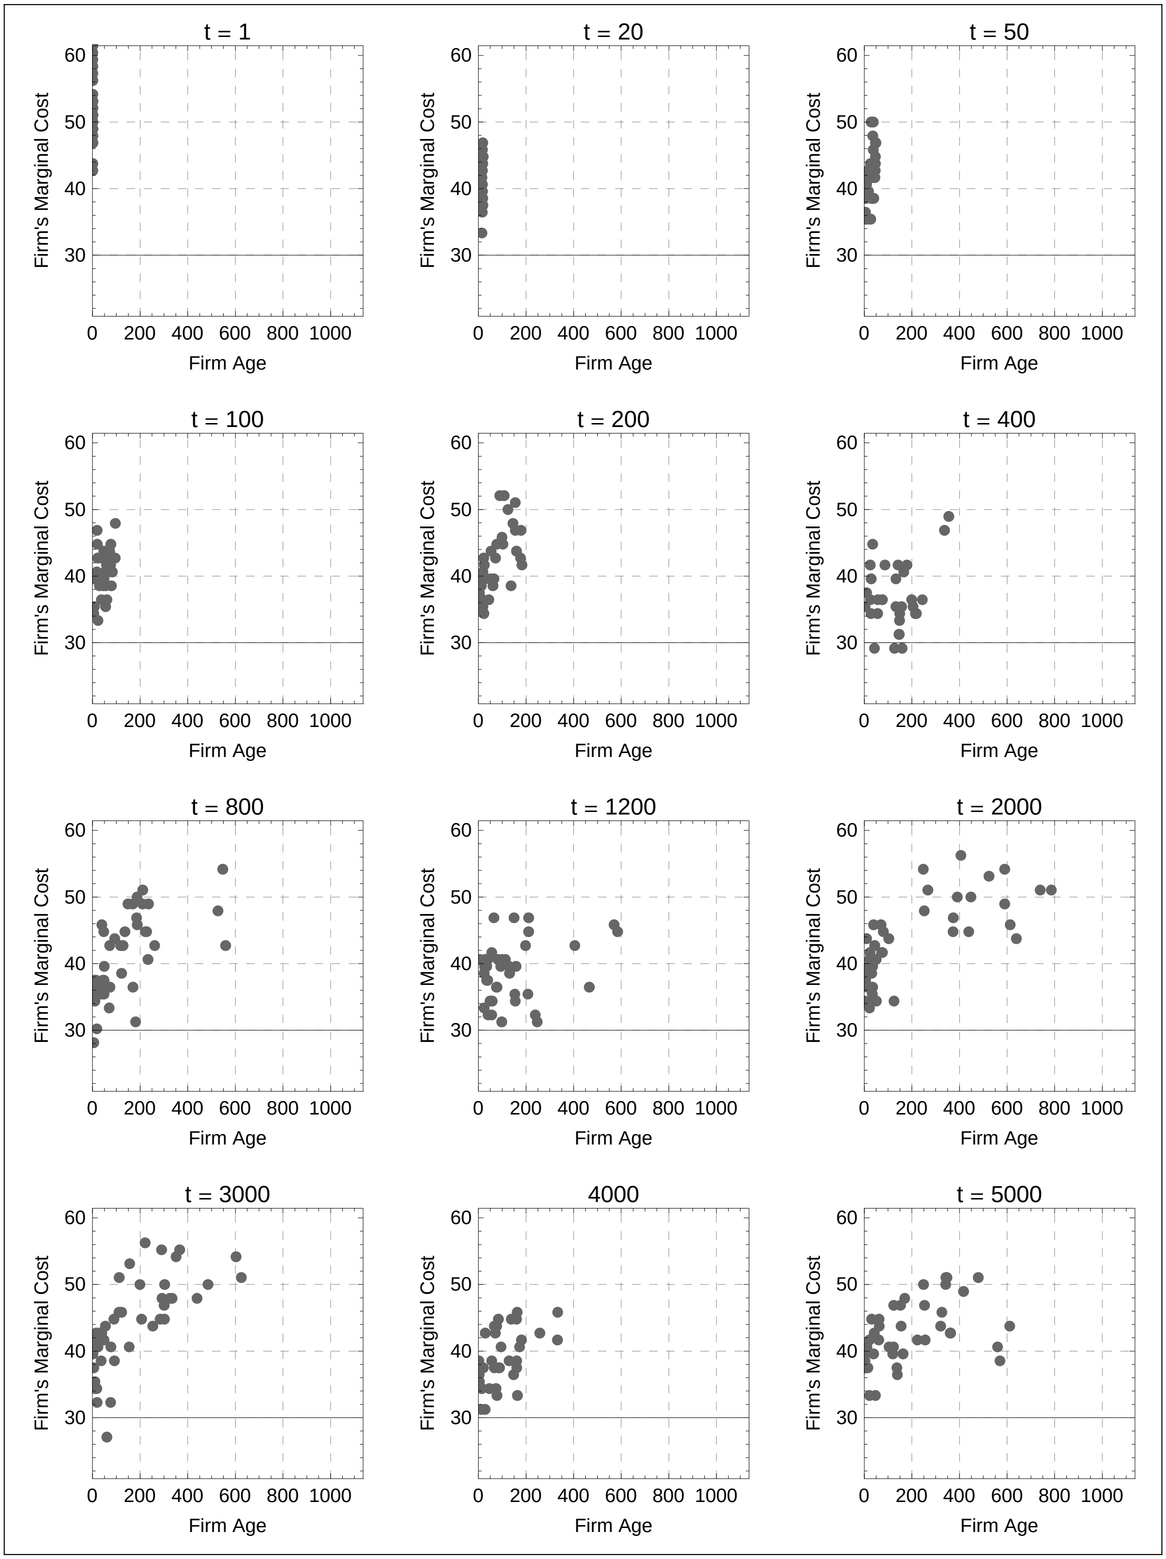

Figure 7: Firm Age and Marginal Costs

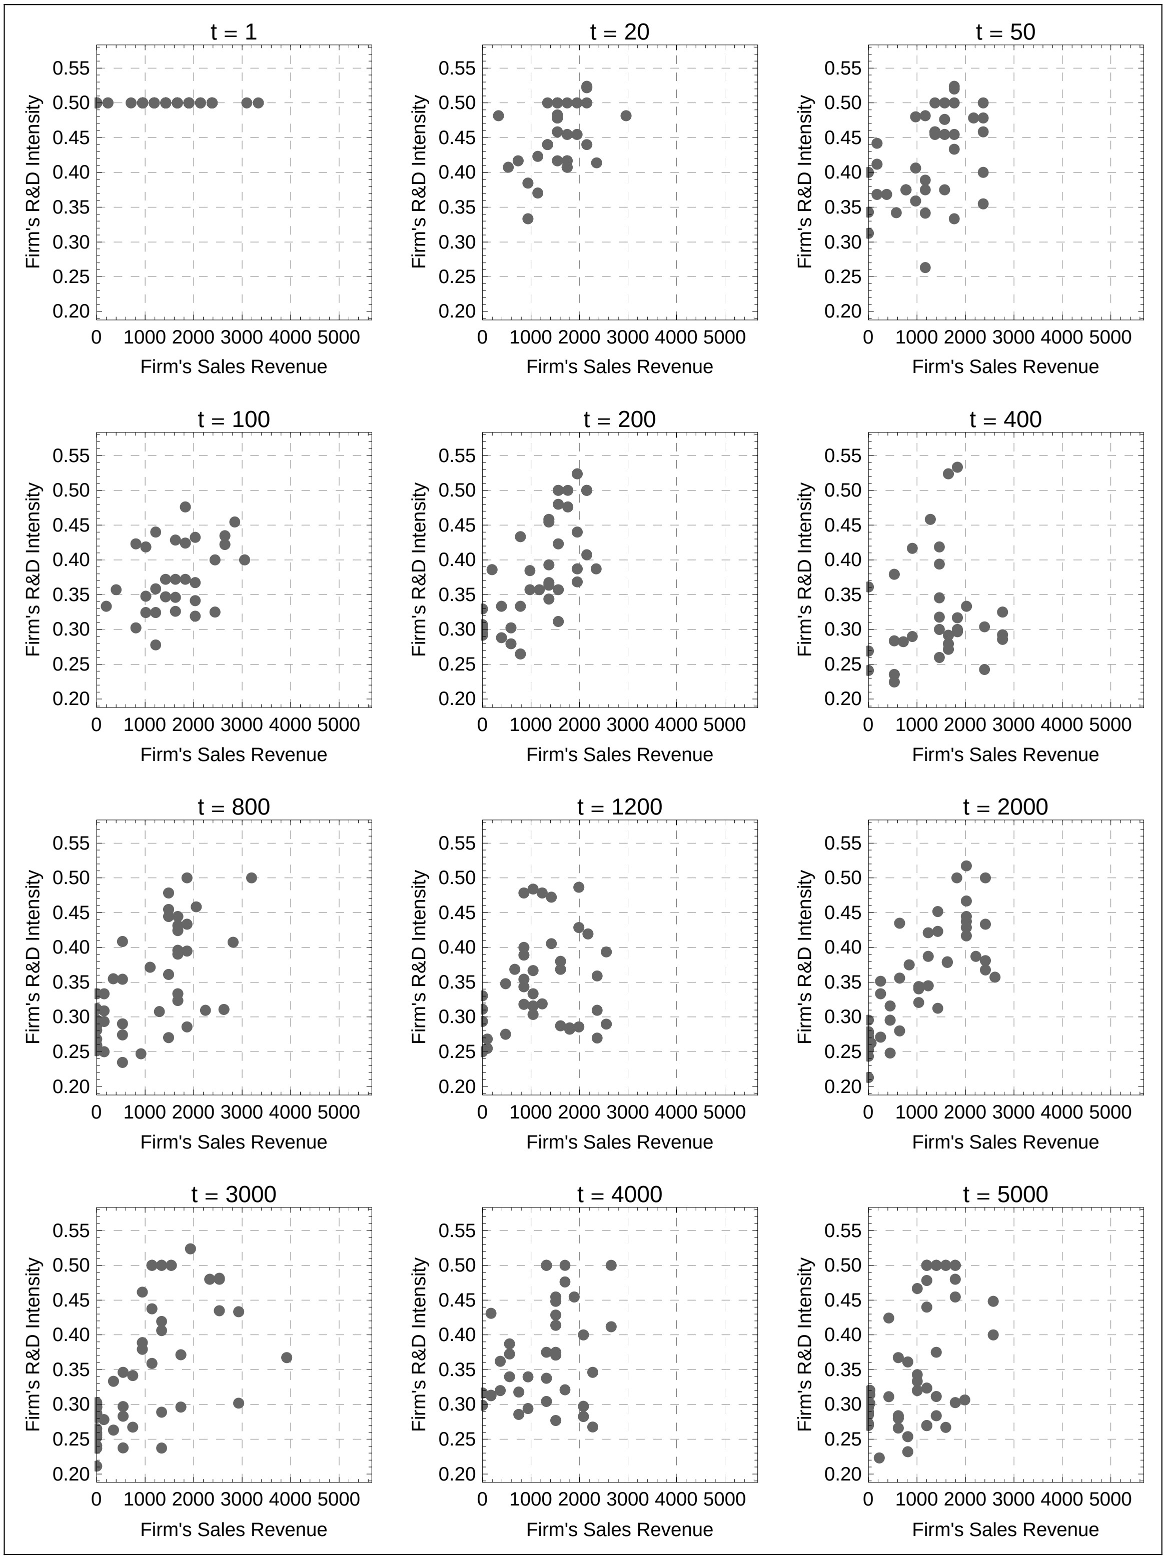

Figure 8: Firm Sales and R&D Intensity

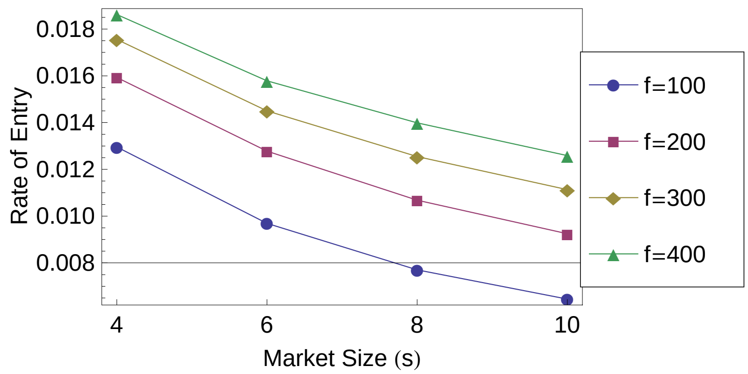

Figure 9: Rate of Entry

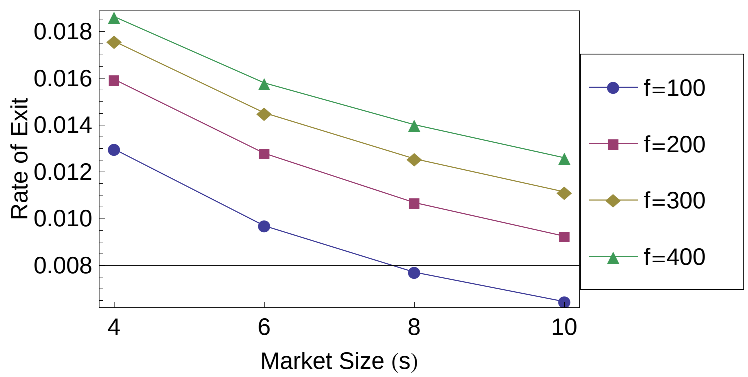

Figure 10: Rate of Exit

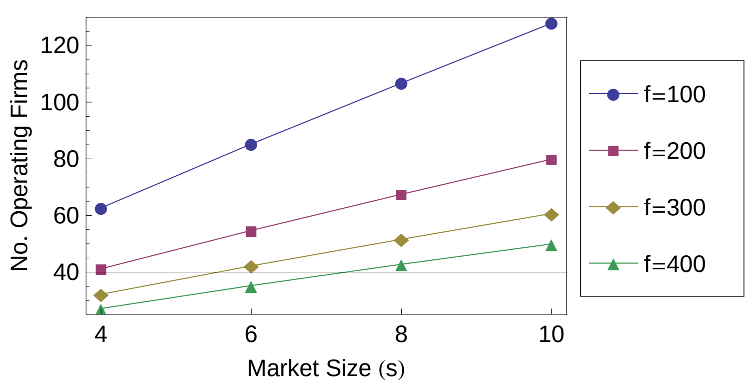

Figure 11: Number of Operating Firms

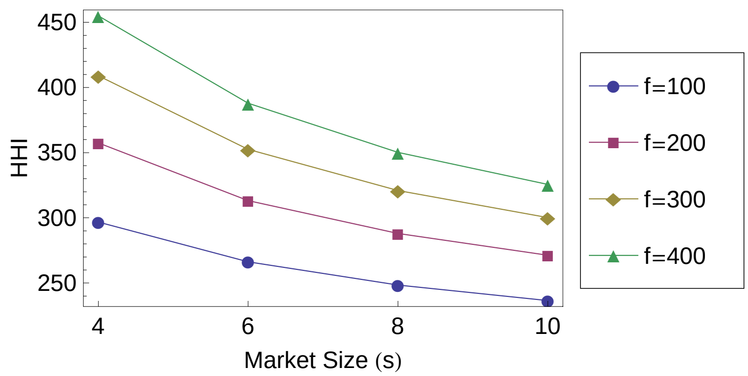

Figure 12: Industry Concentration (HHI)

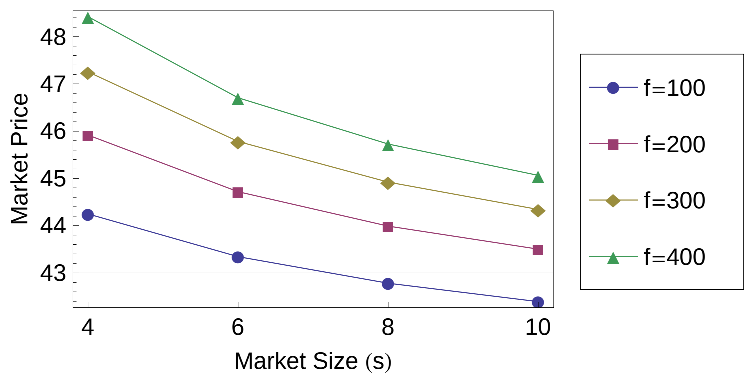

Figure 13: Market Price

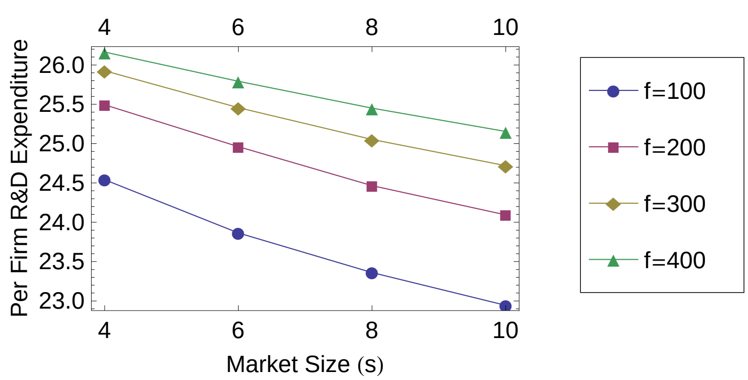

Figure 14: Average R&D Expenditure per Firm

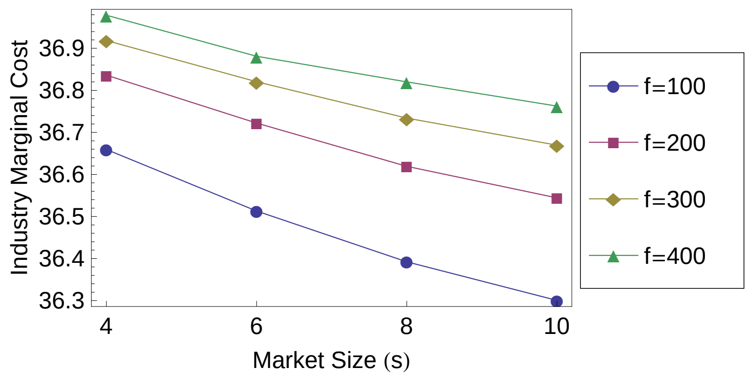

Figure 15: Industry Marginal Cost

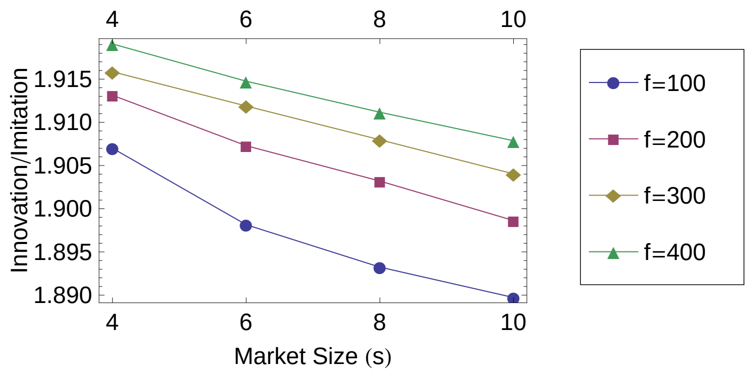

Figure 16: Innovation-to-Imitation Cost Ratio

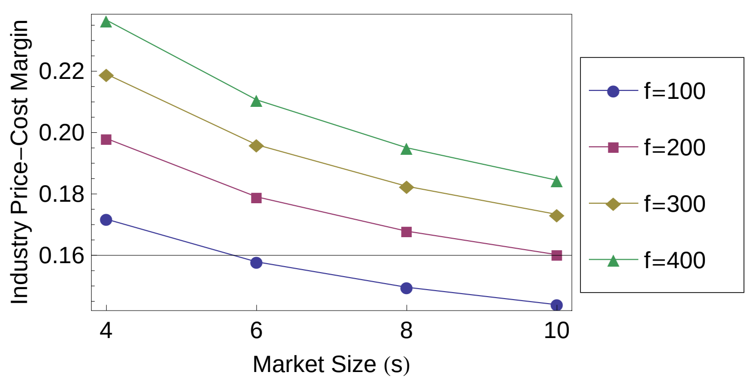

Figure 17: Industry Price-Cost-Margin

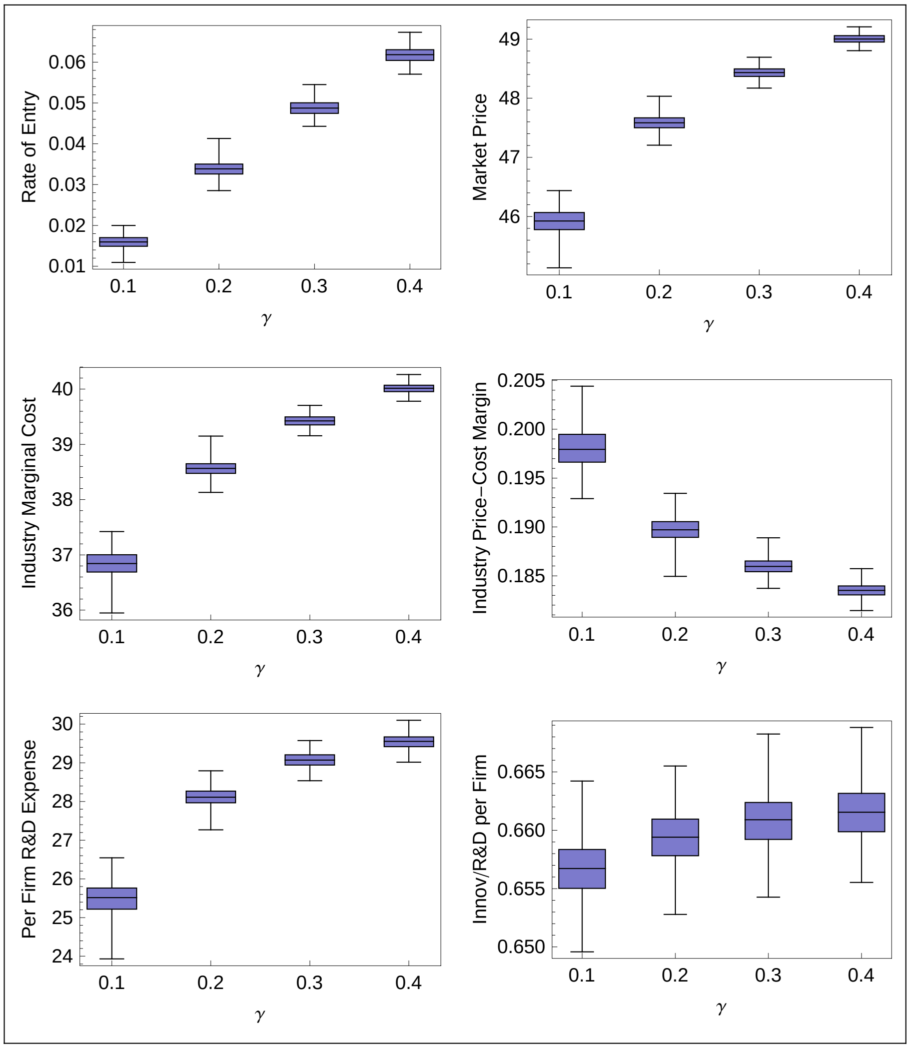

Figure 18: Impact of γ