Industry Dynamics: Aggregate Uncertainty, Heterogeneity, and the Entry and Exit of Firms*†

Mallika Ishwaran‡

June 1999

Abstract

This paper examines the issue of investment under uncertainty. The model developed is a dynamic equilibrium model of a competitive industry. The industry consists of a continuum of heterogeneous firms, each of which faces both aggregate and idiosyncratic uncertainty. The model also allows for the entry and exit of firms in equilibrium. This paper makes two contributions. First, it provides a framework within which to study factors affecting variables such as entry and exit rates and changes in industry size and composition over time. It also studies the impact of differences in decision making at the firm level on aggregate dynamics. The model replicates many of the features observed in the data on firm and industry dynamics. Second, it provides a methodological tool which can be used to solve similar models. In particular, this paper develops computational methods that can be used to characterize the equilibrium behavior of the industry in response to aggregate shocks.

---

*Preliminary and incomplete.

†I would like to thank Tony Smith for all his help and guidance on this project. I would also like to thank Finn Kydland, Steve Spear, Dean Corbae, and participants at the economics seminar series at GSIA for useful comments and insights.

‡Graduate School of Industrial Administration, Carnegie Mellon University

1. Introduction

The purpose of this paper is to develop a dynamic model of firm and industry behavior that can be used to understand the relationship between firm-level decisions, aggregate uncertainty and the business cycle. The model assumes heterogeneous firms, and studies the investment behavior of these firms in response to aggregate shocks and its resulting impact on industry dynamics. Dropping the assumption of a representative firm also allows for the incorporation of firm entry and exit into this model. This enables us to study the changes in industry size and composition over time. The paper builds on the work done by Lucas and Prescott [1971]. In their article, Lucas and Prescott develop a dynamic equilibrium model of investment under uncertainty assuming an industry consisting of a large number of identical firms, each of who face an aggregate shock to demand in each period. The extensions to the Lucas and Prescott framework mentioned above allow for a systematic study of factors affecting industry dynamics.

One of the obvious applications for a model incorporating features such as heterogeneity and entry and exit is in the area of policy. Firm-level data does show a great variability in decision-making across firms. Examples of these differences include simultaneous entry to and exit from an industry and differences in both capital growth rates and market failures rates. The implication of heterogeneity is that firms will typically have differing responses to changes in the aggregate variables. The aggregate response of an industry to any change in the aggregate variables will depend on the features of this distribution of responses. Thus, the impact of any fiscal policy (such as investment tax credits, capital gains tax etc.) will depend on the distribution of responses across all firms in the industry. Another application of this model is in the area of regulatory policy as the distribution of responses is an important factor when studying the impact of regulatory changes on the market structure. Thus, in evaluating the impact of various policy changes (fiscal and regulatory), the characteristics of the distribution (of response heterogeneity) are important and this fact needs to be incorporated into any model aimed at studying these issues.

This paper makes two contributions. First, it develops a model of investment behavior that reflects features observed in microeconomic data on firm behavior. It then studies the impact of incorporating these features of the distribution on aggregate investment dynamics and the evolution of industries. In relaxing assumptions such as that of a representative firm, this paper provides a richer and more realistic framework within which to study investment dynamics. The second contribution of the paper is methodological. In the absence of an analytical solution, it develops a computational technique that allows for such a model to be solved for an approximate (numerical) equilibrium.

Incorporating heterogeneity into the Lucas and Prescott framework is consistent with the empirical evidence on firm investment behavior. Data on firm behavior indicates that the investment decision of any single firm is not perfectly correlated with that of other firms in the industry. This is contrary to the results emerging from representative firm models where all firms make perfectly synchronized investment decisions in response to aggregate fluctuations.

As reported in Doms and Dunne [1998], Troske [1996] and Dunne, Roberts and Samuelson [1989], heterogeneity in firm adjustment decisions is observed both, across time as well as across firms. Looking at the time series, the pattern that emerges is that firms have a few periods of intense capital growth where they adjust their capital stock by at least 37% and many periods of relatively small capital adjustment[1]. These large investment episodes at the firm-level do have an impact on aggregate investment. As the frequency of plants undergoing large investment episodes in any given period goes up, so does aggregate investment[2]. In fact, as documented by Doms and Dunne [1998], aggregate investment tends to be skewed with a small number of firms accounting for a large share of aggregate investment[3]. The cross-sectional data on firm behavior indicates that firms tend to differ in characteristics such as age and size. As firms age, there is a decline in their capital growth rates (see Evans [1987]) and an increase in their relative size[4]. Further, as firms grow in size their market failure rate also declines (see Evans [1987]). All these facts imply that an individual firm's response to aggregate fluctuations is varied and not perfectly correlated with that of other firms, and that individual firms tend to have different capital adjustment paths and growth processes.

This paper also studies the impact of entry and exit on industry equilibrium. Heterogeneity among firms allows for the possibility that some firms will choose to enter or exit the industry in each period. It has been observed empirically that industries tend to evolve over the course of the business cycle with the birth, growth and death of firms. As shown in Hause and Rietz [1984], the entry rate is positively correlated with aggregate fluctuations and the business cycle. Periods of high growth in output also correspond to periods of high entry. Cyclicality in entry rates imply cyclical fluctuations in the exit rates as a large fraction of firms that enter, tend to exit in the following periods (see Dunne, Roberts and Samuelson [1988]). Studying entry and exit in a framework with aggregate uncertainty allows us to replicate the dynamic patterns and correlations observed between entry, exit and the business cycle. It also allows us to make predictions regarding the impact of aggregate uncertainty on cohort behavior[5] and firm-level uncertainty.

One of the important findings of this model is that firm-level heterogeneity does have an impact on the aggregate variables. The measure of concentration of the industry (in terms of investment) is positively correlated with aggregate investment implying that periods of high concentration (with a few firms accounting for a large share of industry investment) are also periods of high aggregate investment. This result is consistent with the empirical

[1] As reported in Dome and Dunne [1998], of the 16 periods for which there is data, in 12 of those periods, average capital growth rate for a firm is between -10% and $+10\%$. In the remaining 4 periods, the average capital growth rate varies from $+12\%$ to $+46\%$.

[2] The time series for frequency of firms undergoing large investment episodes and aggregate investment are positively correlated and tend to move together over the course of the business cycle (see Doms and Dunne [1997]).

[3] On average 25% of investment occurs in firms that are increasing their capital stocks by more than 30%. However, these firms make up only 8% of the sample.

[4] Average firm-size changes over time. Entering firms are small (approximately one-third the size of an average firm in the industry). 20 years later, firms that have survived are approximately one-and-a-half times the size of an average firm (see Doms and Dunne [1998]).

[5] A cohort is defined as a group of firms that enter in the same period.

evidence as reported in Doms and Dunne [1998]. Another interesting implication of the model is that it replicates the cyclical patterns of firm entry and exit observed in the data. Periods of industry growth also correspond to periods of high entry and low exit. Looking at the impact of aggregate uncertainty on the size and composition of entering cohorts, one finds that entering firms tend to be smaller but the cohort itself tends to be larger for good realizations of the aggregate shock. This fact, along with the cyclical behavior of entry and exit, has implications for the change in size and composition of the industry over the course of the business cycle. Thus, in this model, periods of growth are also periods in which industry size is large but individual firms themselves are small.

The methodological contribution of this paper is that it provides a way of computing an approximate equilibrium for a model with the features described above. Assuming a representative firm makes easier the task of computing industry equilibrium as the aggregation of the optimal decision rules of individual firms will exactly give the decision rule for the industry as a whole. Existence of heterogeneity (in the form of a firm-specific productivity shock in each period) implies that the aggregation property will no longer hold. This leads to imperfections in firm predictions of industry price as there is no easy way of aggregating output from the firm-level to the industry-level. In order to be able to forecast future prices precisely, each firm will need to keep track of the distribution of capital (and, hence, output) across firms and how this distribution evolves over time.

As there exists no close-form solution for the model, this paper uses computational methods to solve for industry equilibrium. The solution method is similar to the one developed by Krusell and Smith [1998] and requires firms to keep track of only a few moments of the capital distribution in order to be able to forecast prices accurately. In their paper, Krusell and Smith [1998] look at the impact of income and wealth heterogeneity in a stochastic growth model and find that using one moment (the mean) to approximate the distribution of wealth is adequate and that including further moments has a very small impact on the computed equilibrium. The computational method developed in this paper adapts their technique to the model studied in this paper. It uses two moments (total capital stock and the fraction of firms that decide to enter) and the aggregate shock in order to forecast prices. Using a finite number of moments of the capital distribution to forecast prices reduces the problem from an infinite-dimensional one to a finite-dimensional one. In using this approximation, there is a loss of accuracy but, as shown later, the errors in forecasting prices arising from this approximation are very small.

This paper uses a model similar to the one developed by Hopenhayn [1992] and Hopenhayn and Rogerson [1993] who study firm dynamics in a competitive industry. In their papers, Hopenhayn [1992] and Hopenhayn and Rogerson [1993] construct a model with firm-level heterogeneity but with no aggregate uncertainty. Firms receive an idiosyncratic shock in each period which then leads them to make a decision about whether to exit or stay in the industry. They, then adjust their capital stock accordingly. This framework does provide a valuable insight into firm-level dynamics and the behavior of cohorts over time but it does not look at how the differences in microeconomic adjustment behavior of firms interact with fluctuations in the aggregate variables. By introducing aggregate uncertainty and allowing the aggregate variables (such as investment and output) and the entry and exit rates to fluctuate over time, this paper helps better understand the relationship between firm-level dynamics, aggregate fluctuations and the business cycle.

Other related work include that of Campbell and Fisher [1998] who study the variation in job creation and job destruction rates in the labor market. In their model, entry is assumed to be instantaneous. This approach to modeling entry is useful when thinking of employment creation and destruction decisions but is counter-intuitive when thinking of firms making capital creation and destruction decisions. As noted by Mansfield [1962] and Orr [1974], barriers to entry lower entry and exit rates. In this paper, entry is modeled as a short run phenomenon where firms pick, a period in advance, the optimum level of capital with which to enter the industry. The number of firms that enter the industry in the following period is such that the expected value of entering (across both the aggregate and the idiosyncratic states) is equal to the cost of entering the industry. Instantaneous entry, as in Campbell and Fisher [1998], avoids the aggregation issue as any incumbent firm, after observing the aggregate shock, can predict the industry price exactly by solving for it from a typical entrants' entry condition. For any other case, where there are frictions to the process of entry, firms have to keep track of the distribution of capital in order to be able to forecast industry prices accurately.

Pakes and McGuire [1994] and Ericson and Pakes [1995] develop and solve models of industry dynamics under conditions of aggregate and idiosyncratic uncertainty. Ericson and Pakes solve a Cournot-Nash, homogeneous product model of industry behavior where firms are differentiated on the basis of their efficiency in production. The level of efficiency for any firm is determined by its investment decision and an aggregate process which shifts the cost of the factors of production for the industry as a whole. The algorithm developed to compute the Markov-perfect Nash equilibria (see Pakes and McGuire [1994]) is useful for analyzing industries with a small number of firms but is not adequate for analyzing larger markets. The technique becomes computationally demanding for industries that have a large number of firms active in each period. Thus, it cannot be easily applied to study differences in firm behavior and patterns of evolution for industries consisting of more than a few firms.

The remainder of the paper is organized as follows. The next section explains, in more detail, the model to be used in this paper. It lays out the dynamic framework within which the analysis is conducted and the definition of equilibrium in this industry. Section 2.3 explains the computational strategy used when solving for the equilibrium. Section 2.4 sets out the calibration procedure and parameter values used. Section 2.5 reports the results and the important implications of the model. In Section 2.6, I conclude by laying out the main findings of the paper and the areas for future research in this topic.

2. Model

2.1. Basic Framework

This paper develops a dynamic equilibrium model of industry investment under uncertainty. The model used is that of a competitive industry made up of a continuum of heterogeneous firms. The industry expands and contracts with the entry, growth and exit of firms, implying a change in the structure and size of the industry over time.

2.1.1. Incumbent Firms

We begin this section by describing the decision problem faced by an incumbent firm in the industry. We consider a competitive industry with a continuum of heterogeneous firms, each firm of measure zero. The total measure of firms in the industry fluctuates as firms enter and exit the industry. Each firm produces single good which is identical to that produced by any other firm in the industry. Production of this good involves the use of a single factor of production, capital. Firms differ from each other in terms of the productivity shock that they receive in each period. Output produced by a single firm in any period depends on its existing capital stock and its realization of the productivity shock for that period. Thus, the production technology for any firm in the industry is given by

$$y_{it} = f(k_{it}, z_{it})$$

where f is concave and increasing in both its arguments.
k_{it} is the capital stock of firm i in period t and z_{it} is that firm's realization of the idiosyncratic shock.

The firm specific shock takes on values in R_+ and follows a first-order Markov process described by a function $F(z_{it}, z_{it+1})$ where for each current value of the shock z_{it}, $F(z_{it}, z_{it+1})$ is the distribution function for next period's value of the shock, z_{it+1}. These idiosyncratic shocks are independent across firms (such that the idiosyncratic shocks are not a source of aggregate uncertainty), but, for any individual firm, the shock evolves according to the function F.

The inverse demand curve for the industry is given by

$$p_t = g(Y_t, \theta_t)$$

where Y_t is aggregate output in time t and θ_t is the realization of the aggregate shock in time t.

The aggregate shock also follows a first-order Markov process with a time-invariant transition probability function, G, which determines the value of next period's shock, θ_{t+1}, given the current value of the shock, θ_t. The aggregate shock θ can take on one of two values, i.e., in good states $\theta = \theta_g$ and in bad states $\theta = \theta_b$. A good shock shifts the aggregate demand curve up while a bad shock shifts the aggregate demand curve down. The transition matrix for the aggregate shock is given by

$$\begin{bmatrix} \pi_{ss} & \pi_{s's} \\ \pi_{ss'} & \pi_{s's'} \end{bmatrix}$$

where $\pi_{ss'}$ is the probability that the aggregate shock next period is $\theta_{s'}$ given that it is θ_s this period.

Each firm faces a cost of adjusting its capital stock, i.e.,

$$\text{Adjustment Cost} = h(k_{it}, k_{it-1})$$

where h is assumed to be smooth and convex.
In each period, firms also incur a per-unit cost of using capital, $(r + \delta)$, where r is the opportunity cost of capital and δ is the depreciation rate. This cost is assumed to be determined exogenously. Lastly, each incumbent firm also faces a fixed cost of staying in the industry. The fixed cost, c_f, is introduced in order to make a meaningful distinction between firms exiting the industry and firms temporarily adjusting their capital stock to zero. Thus, c_f can be thought of as a fixed operating cost incurred by the firm in each period that it stays in the industry. This cost is expressed in units of output.

The decision problem of an incumbent firm is as follows. In each period, incumbent firms adjust their capital stocks so as to maximize current and expected value of future profits, taking as given the behavior of all other firms in the industry. Into this optimization problem, they incorporate the fact that they have the option of exiting the industry in the next period. In other words, in period t, after observing both, their idiosyncratic shock and their aggregate shock, incumbent firms decide whether to remain in the industry or to exit in $t + 1$. If they choose to exit, firms implicitly choose tomorrow's capital stock, k_{it+1}, to be zero and thus, incur the adjustment cost $h(0, k_{it})$ but avoid paying the fixed cost, c_f, of staying in the industry in the next period. If a firm chooses to exit, it disappears from the industry, making zero profits in all future periods. If a firm chooses to stay in the industry, it picks the optimal capital level, k_{it}, so as to maximize current and future profits of staying in the industry. This process is repeated in each period.

A key element of a recursive equilibrium for the industry is a law of motion for the aggregate state of the industry. The two aggregate state variables are θ_t and Γ_{t-1} where θ_t is the realization of the aggregate shock in period t and Γ_{t-1} is the period $t - 1$ measure (or distribution) of firms over levels of capital and values of the idiosyncratic productivity shock. The aggregate shock θ evolves according to the exogenous law of motion G and the distribution Γ evolves according to a deterministic law of motion denoted by H such that

$$\Gamma_t = H(\Gamma_{t-1}, \theta_t)$$

Thus, aggregate output in time t is given by

$$Y_t = W(\Gamma_t)$$

where W requires integrating the individual firms' production functions against the distribution Γ_t.

The period t market clearing price can be expressed as a function of aggregate state variables, i.e.,

$$p_t = g(W(\Gamma_t), \theta_t)$$

For an individual firm, the relevant state variables are the holdings of capital with which it enters period t, its realization of the idiosyncratic productivity shock in period t, and the aggregate state, i.e., the relevant state variables are $(k_{it-1}, z_{it}, \Gamma_{t-1}, \theta_t)$. The role of the aggregate state is to allow the firms to compute current and future industry prices.

The optimization problem faced by an incumbent firm in each period can therefore, be expressed by the following Bellman equation:

$$\begin{aligned} V(k_{it-1}, z_{it}; \Gamma_{t-1}, \theta_t) &= \max_{k_{it}} \{p_t y_{it} - (r+\delta)k_{it} - h(k_{it}, k_{it-1}) - p_t c_f + \\ &\quad \beta \max E[V(k_{it}, z_{it+1}; \Gamma_t, \theta_{t+1}) | z_{it}, \theta_t], -h(0, k_{it})]\} \end{aligned} \tag{2.1}$$

subject to

$$\begin{aligned} y_{it} &= f(k_{it}, z_{it}) \\ Y_t &= W(\Gamma_t) \\ p_t &= g(W(\Gamma_t), \theta_t) \\ \Gamma_t &= H(\Gamma_{t-1}, \theta_t) \end{aligned}$$

and the stochastic laws of motion for θ and z as given by G and F, respectively.

The nested max operator in the Bellman equation reflects the fact that firms implicitly incorporate their exit decision into their optimization problem. In picking the optimal level of capital in time t, firms evaluate the expected value of remaining in the industry (i.e., the expected value of future profits) and compare this with the present discounted value of profits associated with exiting the industry the following period (i.e., $-h(0, k_{it})$). Thus, the investment decision in period t is dependent on the decision to exit in period $t+1$.

The solution to this optimization problem is the decision rule for capital

$$k_{it} = s(k_{it-1}, z_{it}; \Gamma_{t-1}, \theta_t)$$

The law of motion for the evolution of the industry

$$\Gamma_t = H(\Gamma_{t-1}, \theta_t)$$

can be derived by summing across all firms, each of who use s to make their investment decisions.

For the optimization problem given by (2.1), a threshold rule characterizes the incumbent firm's exit decision, i.e., there exists a $k^* = k_{it-1}$ such that

$$E[V(k_{it}, z_{it+1}; \Gamma_t, \theta_{t+1}) | z_{it}, \theta_t] = -h(0, k_{it})$$

where k_{it} is the solution to the optimization problem (2.1) and is determined by s.
For any $k_{it-1} < k^*$,

$$E[V(k_{it}, z_{it+1}; \Gamma_t, \theta_{t+1}) | z_{it}, \theta_t] < -h(0, k_{it})$$

and the firm will choose to exit. For all other values of k_{it-1}, the expected value of remaining in the industry is greater than that of exiting and the firms will choose to stay. Thus, the threshold value for k is defined as

$$x(z_{it}, \Gamma_{t-1,} \theta_t) = k^*$$

such that

$$E[V(k_{it}, z_{it+1}; \Gamma_t, \theta_{t+1})|z_{it}, \theta_t] = -h(0, k_{it}) \tag{2.2}$$

where, given the optimum decision rule s, k^* is the level of capital such that , for any individual firm, the value of staying in the industry is equal to that of exiting in the next period.

2.1.2. Entering Firms

Entry is modeled such that firms make their entry decision one period in advance of when they actually enter the industry. All firms are ex-ante identical, i.e., they enter with the same amount of capital stock and expect to receive their idiosyncratic shock from some invariant distribution that is common to all firms. Upon entering the industry, each new entrant receives its draw of z from a time-invariant distribution ν which is independent of the number of entering firms. These draws are independent and identically distributed across all entering firms. The level of capital stock, k_t^e, with which firms decide to enter is determined by entering firms maximizing the expected value of entering, i.e.,

$$\max_{k_t^e} E[V(k_t^e, \nu_{t+1}, \Gamma_t, \theta_{t+1})|\nu_t, \theta_t]$$

subject to

$$\Gamma_t = H(\Gamma_{t-1}, \theta_t)$$

and the law of motion for θ as given by the transition probability matrix, G, and the invariant distribution of the z-shock, ν.

The solution to the above optimization problem is the level of capital with which firms choose to enter the industry in $t+1$, i.e.,

$$k_t^{e*} = S(\Gamma_{t-1}, \theta_t) \tag{2.3}$$

The expected value of entering this industry is given by

$$E[V(k_t^{e*}, \nu_{t+1}, \Gamma_t, \theta_{t+1})|\nu_t, \theta_t]$$

Each entering firm incurs a one-time entry cost, c_e. As in the case of c_f, the entry cost is denominated in units of output. As entry is unrestricted (in the sense that there is no limit on the number of firms that can enter the industry in any period), the expected value of future profits of an entering firm must be less than or equal to the cost of entry, i.e.,

$$E[V(k_t^{e*}, \nu_{t+1}, \Gamma_t, \theta_{t+1})|\nu_t, \theta_t] \leq p_t c_e \tag{2.4}$$

Thus, the entry condition is such that there is no incentive for more firms to enter.

2.2. Definition of Equilibrium

In order to define a competitive equilibrium, it remains to define the equilibrium entry condition, the aggregate quantities and the transition rule for the $\Gamma-$distribution. For an equilibrium with positive entry, it must be that the above condition (2.4) holds with equality, i.e., the number of firms entering should be such that the value of entry, in expectation, is equal to the cost of entering. Thus, the entry rule

$$E[V(k_t^{e*}, \nu_{t+1}, \Gamma_t, \theta_{t+1})|\nu_t, \theta_t] = p_t c_e \tag{2.5}$$

must hold in each period.
Total industry capital is

$$K_t = \int k_{it} d\Gamma_t(k_{it}, z_{it}) \tag{2.6}$$

and total industry output is

$$Y_t = \int y_{it} d\Gamma_t(k_{it}, z_{it}) \tag{2.7}$$

where Γ_t is the distribution of firms across different levels of capital and values of the z-shock in time t.

A competitive equilibrium is reached if the following conditions are satisfied:

1. The market clears, i.e.,
$$p_t = g(W(\Gamma_t), \theta_t)$$
2. There exists a unique solution to equation (2.1) given the stochastic processes for the aggregate state variables.
3. The function s solves the incumbent firms' optimization problem.
4. $x(z_{it}, \Gamma_{t-1}, \theta_t)$ is defined as in (2.2).
5. The entry condition (2.4) holds with equality in every possible state with positive entry.
6. Aggregate capital, K_t, is defined as in (2.6).
7. Aggregate output, Y_t, is defined as in (2.7).
8. The law of motion for the evolution of the distribution of firms, H, is generated by the firms' optimal decision rule, s.

If the above conditions are satisfied, the industry can be said to be in competitive equilibrium.

3. Description of the Computational Algorithm

In the absence of any close-form solution for the model described above, numerical methods are used to compute an approximate equilibrium. As the model assumes a continuum of heterogeneous firms, numerically solving for industry equilibrium is very expensive from the computational point of view. The reason for this is that one of the aggregate state variables for this industry is the distribution of firms (Γ) across holdings of capital and values of the idiosyncratic productivity shock. Thus, firms have to keep track of how this distribution evolves over time in order to forecast future industry prices. The approach developed in this paper is similar to that used by Krusell and Smith [1998] for the analysis of a general equilibrium macroeconomic model with income and wealth heterogeneity. This approach requires firms to use a finite set of moments of the $\Gamma-$distribution in order to be able to forecast future prices. The computational burden is thereby reduced as we now have to keep track of a finite set of state variables and not the entire distribution. The errors in forecasting (of industry price) arising from the use of this method are extremely small (as reported in Section 2.5) and, in this sense, the approximate equilibrium so computed is close to the actual one[6].

The basic idea underlying this numerical strategy is to use a finite set of the first I moments of the $\Gamma-$distribution, when forecasting future industry prices. Let m_t be a vector of the first I moments of Γ_t, i.e.,

$$m_t = (m_{1t}, m_{2t}, ..., m_{It})^7$$

In order to approximate the law of motion for the distribution, this computational algorithm uses a class of functions H_I which express m_{t+1} (the vector of I moments in the next period) as a function of the current I moments, i.e.,

$$m_{t+1} = H_I(m_t, z_t)$$

The second approximation in this procedure is that of the function W, which maps the $\Gamma-$distribution into aggregate output Y, that is,

$$Y_t = W_I(m_{1t}, m_{2t}, ...m_{It})$$

Assuming H_I and W_I to be given, a typical firm then solves its optimization problem. The resulting optimal decision rule for capital is s_I. We use this decision rule to generate time series data for a large number of incumbent firms, i.e., we generate a time series for the above mentioned I moments of the distribution.

[6] The phrase "approximate equilibrium" refers to the fact that firms do not include all the moments of the $\Gamma-$distribution (and hence do not use all the information available to them) when computing current and future industry prices.

[7] In this case, m_{1t} is a 2×1 vector consisting of the first moments of Γ_t (i.e., the average capital holdings and the average value of the idiosyncratic shock), m_{2t} is a 2×2 matrix consisting of the second moments of Γ_t (i.e., the variance of capital holdings across firms, the variance of the idiosyncratic shock across firms and the covariance between capital holdings and the idiosyncratic shock), and so on.

In each period, the number of firms that will enter the industry in the following period is determined by imposing the entry condition (2.5). The level of capital stock that entrants will enter with is determined by (2.3). Once firms have entered the industry, their investment decisions are made using s_I.

The approximation to the industry equilibrium is then a pair of functions H_I and W_I which, when taken as given by the firms, satisfy the two following conditions:

- Each function yields the best fit within its class to the resulting simulated time series.
- Each function yields a fit that is close to perfect. Specifically, H_I tracks the evolution of m_t in the simulated time series almost exactly and W_I predicts Y_t as a near exact function of m_t.

Thus, in the computed approximate equilibrium, firms do not take into account all the moments of the distribution when solving their optimization problem (specifically, they don't take into account all the moments of the distribution when computing current and future industry prices) but the errors in forecasting that result from this omission are very small.

Elaborating further, the iterative procedure used in this algorithm can be summarized as follows:

1. Choose the number of moments I of distribution Γ that are to be used to forecast industry prices.
2. Assume functional forms for H_I and W_I and guess on the parameters for that functional form.
3. Given H_I and W_I, firms then solve their optimization problem.
4. Use the resulting decision rule to simulate time series data for all incumbent firms over T periods. The decision rule incorporates firms' exit decision. Thus, in each period, a certain number of firms (determined by the optimum decision rule, x) exit the industry. Impose conditions (2.3) and (2.5) in order to determine the number of entering firms and the capital stock of each entering firm for the following period. Once firms have entered the industry, they evolve like any other incumbent firm in the industry and the decision rule derived in Step 3 is used to simulate time series for them.
5. The stationary region of this time series data is used to form new estimates for the parameters of the H_I and W_I functions. These new laws of motion are then used to generate a new or modified decision rule for capital for individual firms. Continued iteration on the function parameters results in a fixed point.
6. At this stage, if the goodness-of-fit of the estimated parameters is good, then an approximate equilibrium has been reached that is close to the actual one. If the goodness-of-fit is not satisfactory, then the options available include repeating the

above process using a larger number of moments and/or using a different functional form for H_I and W_I[8].

The advantage of this algorithm is that a firm has to keep track of a finite set of state variables (as opposed to an infinite set) in order to be able to forecast future prices to a reasonable degree of accuracy. This approximation lowers the computational burden and reduces the problem to manageable proportions. Thus, frugality in the choice of moments and the choice of moments itself becomes important. For example, if firms keep track of only two moments, it may be more important (in terms of forecasting accuracy) to keep track of the mean capital holdings of incumbent firms and entering firms separately rather than, say, the mean and variance of capital holdings across all types of firms. In addition, it is important to emphasize that the algorithm developed produces a good approximation only to the stationary equilibrium of the industry. As noted by Krusell and Smith [1998] and Caballero [1992], the idiosyncratic and aggregate shocks interact in such a way so as to limit the number and types of distributions that can arise in a stationary equilibrium, i.e., the structure of the model puts strong restrictions on the ergodic set of distributions. The numerical algorithm exploits this fact to achieve frugality in the set of moments used to produce a good approximation to the law of motion for Γ.

4. Baseline Model

4.1. Model parameters

The procedure used to calibrate this model and assign parameter values is similar to the one developed by Kydland and Prescott [1982]. Many of the parameter values calibrated are dependent on the length of the time period, which is set to one year. Hopenhayn and Rogerson [1993] use time periods of length five years but their model abstracts away from cyclical fluctuations. As this paper aims at studying the interaction between cyclical fluctuations and firm adjustment decisions, it is necessary to use shorter time periods[9].

The production technology is assumed to be Cobb-Douglas, i.e.,

$$y_{it} = e^{z_{it}} k_{it}^{\alpha}$$

The technology is decreasing returns to scale and $\alpha = 0.36$.

The aggregate shock θ is assumed to be persistent and is calibrated such that the average duration of a good or bad shock is 4 periods, i.e., the transition probability matrix for the aggregate shock is

$$\begin{bmatrix} 0.75 & 0.25 \\ 0.25 & 0.75 \end{bmatrix}$$

[8]It should be noted that, in the simulations, the markets clear exactly, i.e., the approximation W_I is not used to compute the market clearing price. Instead, the algorithm simply adds up the output of the industry's firms and inserts this value into the industry demand curve.

[9]Parameters are calibrated from the Census of Manufactures (which uses five year time periods) but adjusted to reflect annual levels.

The aggregate shock is assumed to take on one of two values, specifically 0.05 or -0.05. This, along with the demand function parameters are calibrated such that industry output fluctuates by 8%. The firm-specific idiosyncratic shocks have a persistence parameter of 0.95 and a variance of 0.04. We use $\beta = 0.95$ and per-unit cost of capital, $(r + \delta) = 0.15$. Adjustment costs are assumed to be smooth and convex and given by

$$h(k_{it}, k_{it-1}) = \frac{d}{2}(k_{it} - k_{it-1})^2$$

The parameters d, c_f and c_e along with the invariant distribution ν are calibrated such that 9% of firms, on average, enter (or exit) the industry in each period. This is the average percentage of firms entering (and exiting) annually as estimated by Dunne, Roberts and Samuelson [1988]. They study at four-digit U.S. manufacturing industries over the period $1963 - 1982$ when calculating the entry and exit rates. The data set used is constructed from plant-level data collected by the Census of Manufactures.

4.2. Solution and Simulation Parameters

The individual firm's dynamic programming problem is solved by computing an approximation to the value function on a grid of points in the state space. The resulting optimal levels of capital are not restricted to being on the grid and are computed using cubic spline and polynomial interpolation techniques. Appendix A contains a detailed description of the algorithm. The simulations are done for $10,000$ periods. In each period of the simulations, incumbent firms evolve according to the optimal decision rule. The number and size of entering firms is determined a period in advance. Once firms have entered the industry, they evolve like any other incumbent firm in the industry. Thus, the industry size varies with the number of firms choosing to enter and exit the industry over the course of the simulations. In estimating the coefficients of the laws of motion, the first $1,000$ periods of the simulations are discarded.

5. Results

5.1. How Good is the Approximation

The computational technique described in Section 3 allows us to compute an approximate equilibrium with very small forecasting errors. The number of moments required to give a good approximation to the equilibrium are two, namely, the total capital stock in the industry (K_{t-1}) and the fraction of firms that make the entry decision in that period (N_{t-1})[10]. Thus, only two moments (K_{t-1}, N_{t-1}) of the distribution Γ_{t-1} are used to forecast industry prices. Using a linear functional form and $(K_{t-1}, N_{t-1}, \theta_t)$ as state variables, the approximate equilibrium is given by:

[10] Actual entry occurs in the following period.

For $\theta_t = \theta_g$, the law of motion for aggregate capital can be written as

$$\begin{aligned} K_t &= a_{0g} + a_{1g} K_{t-1} + a_{2g} N_{t-1} \\ R^2 &= 0.995 \\ \sigma &= 0.00072 \end{aligned}$$

For $\theta_t = \theta_b$, the law of motion for aggregate capital can be written as

$$\begin{aligned} K_t &= a_{0b} + a_{1b} K_{t-1} + a_{2b} N_{t-1} \\ R^2 &= 0.995 \\ \sigma &= 0.00104 \end{aligned}$$

where R^2 is the goodness-of-fit for the regression and σ is the standard deviation of the regression error.

The laws of motion for the second moment, N, are determined in a slightly different way than that for aggregate capital. The entry condition is made to hold with equality in each period of the simulations. Imposing the entry condition allows us to solve for N as a function of last period's aggregate state.

For $\theta_t = \theta_g$, the law of motion for the fraction of firms entering can be written as

$$\begin{aligned} N_t &= b_{0g} + b_{1g} K_{t-1} + a_{2g} N_{t-1} \\ R^2 &= 0.999 \\ \sigma &= 0.0001 \end{aligned}$$

For $\theta_t = \theta_b$, the law of motion for the fraction of firms entering can be written as

$$\begin{aligned} N_t &= b_{0b} + b_{1b} K_{t-1} + b_{2b} N_{t-1} \\ R^2 &= 0.999 \\ \sigma &= 0.0001 \end{aligned}$$

where R^2 is the goodness-of-fit for the regression and σ is the standard deviation of the regression error.

The W_I function which maps the distribution Γ_t into Y_t is given by

$$\begin{aligned} Y_t &= c_0 + c_1 K_t \\ R^2 &= 0.953 \\ \sigma &= 0.005 \end{aligned}$$

Thus, as can be seen from these results, the model performs well in terms of

1. the number of moments chosen to approximate the distribution of capital across firms when forecasting prices and,

2. the functional forms chosen to approximate the H_I and W_I functions.

In their paper, Krusell and Smith [1998] found that, when solving for equilibrium, only the first moment of the distribution of capital across agents (i.e., the mean capital stock held by all agents) seemed to matter when forecasting prices. Incorporating other moments into the computation of equilibrium did improve the fit, but their quantitative importance was very small. This result (of only the mean mattering) hinged on two factors. First, for higher levels of capital stock, agents' optimal decision rule is close to linear. Second, the equilibrium being considered is a stationary one and this set limits on the types of distributions that can occur[11]. The computational methods used to solve for equilibrium in this paper build on those employed by Krusell and Smith.

For the model used in this paper, the first moment is not sufficient to be able to forecast prices with any degree of accuracy. Looking at the decision rules for individual firms, we find that they are close to linear for higher levels of capital stock. Further, as long as the capital stock is large enough, the slopes of the decision rules don't change significantly for different realizations of the z-shock (see Figure 1).

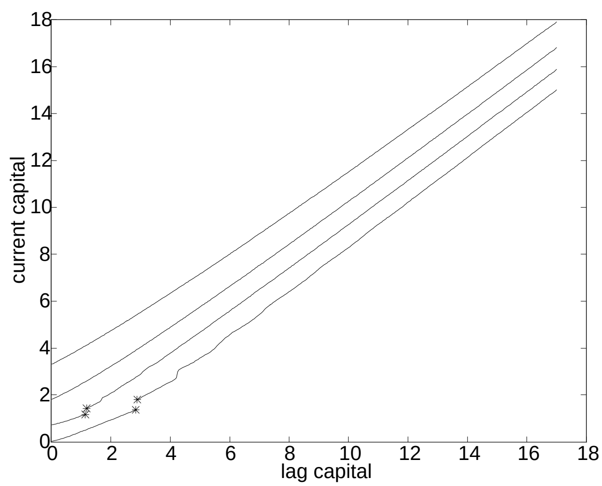

Figure 1: Decision rules for different z-shocks given θ

For lower levels of capital stock, the decision rules corresponding to some realizations of the z-shock have a point of discontinuity (as shown by the two lowest decision rules in Figure 1)[12]. This corresponds to the exit rule described in Section 2. For points to the left of the kink, firms will choose to exit. As the value of the z-shock gets larger, the exit zone becomes smaller until for good enough realizations of the shock, firms choose never to exit (corresponding to the two top decision rules in Figure 1). As can be seen from Figure 1,

[11] For a more detailed discussion of the implications of a stationary equilibrium, see Ericson and Pakes [1995].

[12] The irregularities in the decision rules just to the right of the kink point arise from the fact that the z-shock is made to take on discrete values. For continuous z-values, the decision rules (to the right of the kink point) are smooth.

the threshold capital level, below which firms will decide to exit the industry, is relatively small[13].

Thus, the decision rules for individual firms have significantly different slopes (and, hence, significantly different propensities to invest) only for relatively small levels of capital stock. For higher values of capital, the decision rule is close to linear with similar slopes across different z-values implying that firms with large capital stocks have the same marginal propensity to invest. For firms carrying low levels of capital, the slope of the decision rule is smaller implying a lower marginal propensity to invest. If all firms have the same marginal propensity to invest out of current capital, the redistribution of capital stock will have no impact on aggregate investment. But, if the redistribution occurs in such a way that a large share of the capital stock is redistributed to firms in the lowest range, this will result in large changes in aggregate investment.

The result in Krusell and Smith [1998] will hold true in this case if the stationary equilibrium restricted the types of distributions such that firms tend not to carry low levels of capital stock. This is not the case for the model presented in this paper. In equilibrium, firms are allowed to (and do) enter and exit the industry in each period. Entering firms tend to be much smaller (in terms of capital stock that they enter with) than existing firms in the industry. Exiting firms also tend to be smaller than existing firms in the industry. Thus, in any period, a significant number of firms in the industry have low levels of capital stock. Also, these small firms account for a significant amount of the total capital in the industry[14].

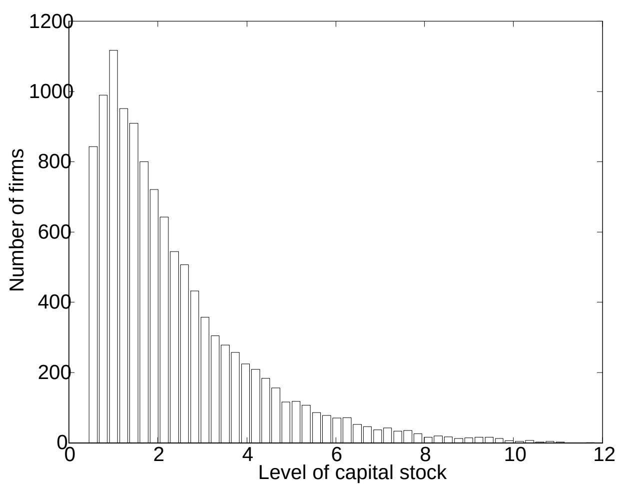

Figure 2: A typical distribution of capital across firms

As seen in Figure 2, a large fraction of firms lie in the lower tail of the distribution[15].

[13]The exit value for k_{it-1} (i.e., values of k_{it-1} below which firms will choose to exit) ranges from 1.19 to 4.25 across different z-values and aggregate states with an average of 2.5.

[14]34% of aggregate capital stock is in the hands of firms who have capital stock of less than 2.5.

[15]On average, 65% percent of firms in any period have capital stocks of less than 2.5.

Table 5.1: **Characteristics of Entering Cohorts**

	Empirical	**Model**
Fraction of Entrants	8%	9.6%
Entrant Relative Size	35.2%	15%
Entrant Market Share	3%	2%

Empirical evidence on the characteristics of entering cohorts is drawn from Dunne, Roberts and Samuelson [1988]. The first entry measures the total number of new firms entering the industry in a given period as a fraction of the total number of firms in the industry in that period. Size of an entrant is measured in terms of the average output produced by an entering firm in the year of its entry, relative to the average output of a firm already in the industry. Market share of a cohort is measured in terms of the cohort's share of aggregate output in the year of entry.

Thus, in an equilibrium with entry and exit, the redistribution of capital is such that there are a lot of firms with low levels of capital and this has a significant impact on aggregate investment. Hence, the result of only the mean mattering does not hold for this model. In order to solve for industry equilibrium, two moments of the capital distribution, K_{t-1} and N_{t-1}, are used by firms when forecasting prices. As entry accounts for a large fraction of firms lying in the lower tail of the distribution, using the fraction of firms entering in any period as a second moment captures this aspect of the distribution. As reported above, two moments of the $\Gamma-$distribution are sufficient for the purpose of forecasting prices and the equilibrium so computed is a good approximation to the actual equilibrium.

5.2. Empirical Implications of the Model

We find that this model provides us with a more realistic framework within which to study investment dynamics. The model performs very well in terms of matching observed data.

5.2.1. Cohort Behavior

The model replicates cohort behavior as has been documented in the data. In their paper, Dunne, Roberts and Samuelson [1988] study the patterns of firm entry, growth and exit for four-digit U.S. manufacturing industries over a 19-year period. They note that, across industries, entrants tend to be much smaller on average than existing firms and account for a very small share of aggregate output. Entrants also have a much higher market failure rate than existing firms with a large number of them exiting within five years of entry. Over time, as cohort-firms age their market failure rate declines as does their capital growth rate while their average size increases (see Table 5.1).

In this model, entering firms tend to be much smaller than the average firm in the industry, i.e., they enter with 11% of the capital stock of an average firm in the industry and produce 15% of the output of an average firm in the industry. In the year of their entry, entrants account for only 2% of total industry output but for approximately 9% of total firms in the industry. They are also prone to higher market failure rates than the average firm in the industry. Almost 79% of a cohort exits within 5 years of entry and 82% exit

Table 5.2: **Evolution of Cohorts Over Time: Model Implications**

	1	**2**	**3**	**4**	**5**	**6**	**7**
Market Share	0.016	0.013	0.015	0.016	0.014	0.013	0.012
Average Size of Surviving Firms	0.15	0.62	0.87	1.11	1.18	1.19	1.193
Market Failure Rate	–	0.7843	0.0274	0.034	0.0207	0.012	0.012
Capital Growth Rate	–	391%	39%	28.5%	15.3%	5%	2.3%

This table looks at the evolution of a cohort over time as implied by the model. It studies the change in cohort characteristics over seven 5-year periods. The first entry measures the cohort's share of aggregate output over time. The second entry looks at the average size of a surviving firm measured in terms of the average output produced by these firms relative to the average output produced by a firm in the industry. The third entry measures the growth in the average capital stock of a surviving firm in the cohort. The last entry quantifies the market failure rate of an entering cohort. It is measured as the fraction of the entering cohort that exit in each period.

Table 5.3: **Evolution of Cohorts Over Time: Empirical Evidence**

	1967	**1972**	**1977**	**1982**
Market Share	0.0278	0.0166	0.0134	0.0106
Average Size of Surviving Firms	0.352	0.597	0.915	1.32
Market Failure Rate	–	0.639	0.151	0.086

This table is drawn from Dunne, Roberts and Samuelson [1988], and it looks at the change in cohort characteristics over time for a cohort that entered in 1967. These figures are the means across 387 four-digit industries for each time period. The change in cohort characteristics is analyzed with respect to the change in market share, average size and the market failure rate of the cohort. The market share is measured as the cohort's share in aggregate output. The average size of a surviving firm is measured in terms of output and is relative to the average size of all firms present in the industry. For example, firms entering in 1967 produced 35.2% of the output of an average firm in the industry. The market failure rate is measured in terms of the fraction of the entering cohort that exits in each period.

within 10 years of entry. The cohort-firms that survive in the industry get much larger in size. At their largest, they produce approximately one-and-a-half times the output of an average firm and have twice the amount of capital stock of an average firm in the industry. Despite getting larger, the capital growth rate of firms tends to decline over time. Firms that do not exit have very high capital growth rates in the first few years following entry after which the growth rate declines rapidly. Thus, the market failure rate and capital growth rate of firms declines over time while the average size of the firm increases. These results are summarized in Table 5.2.

For a cohort as a whole, as reported in Dunne, Roberts and Samuelson [1988], we find that it is responsible for its largest share of aggregate output in the first few years after entry. This decline in the market share of a cohort over time is the result of two conflicting forces, namely, the increase in the size of surviving members of the cohort and the high exit rate of firms in the cohort. Empirical evidence shows that the decline in the number of firms in a cohort more than offsets the increase in average size of surviving firms, leading to a falling market share (see Table 5.3).

Table 5.4: **Characteristics of Entering Cohorts Conditional on the Aggregate Shock**

Fraction of Entrants	1.19
Entrant Relative Size	0.995
Entrant Market Share	1.2
Entering Capital Stock	0.995
Share of Aggregate Capital	1.196

This table quantifies the characteristics of a cohort that enters during good realizations of the aggregate shock, relative to those that enter during bad realizations of the aggregate shock. For example, the fraction of firms entering in good times is 19% more than the fraction that enter in bad times. Each entry in the table is a ratio looking at that cohort characteristic for cohorts entering in good times relative to the same cohort characteristics of cohorts entering in bad times. These numbers are for cohort characteristics in the year of their entry.

As can be seen in Figure 3, the model is able to replicate this feature of the data. We find that entrants' market share (measured by their share of total output), tends to decline over time. The cohort accounts for the highest share of industry output in the first few years after entry. Subsequently, the high market failure rate outweighs the growth in size of the surviving firms, leading to a fall in the cohort's market share.

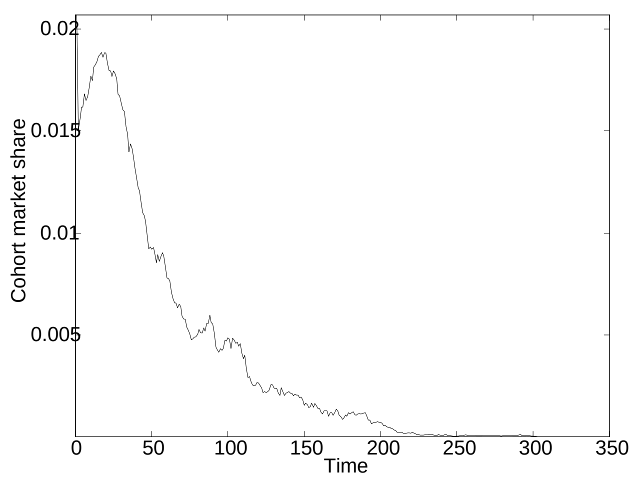

Figure 3: Market share of a cohort over time

5.2.2. Aggregate Uncertainty and Cohort Behavior

Aggregate uncertainty has a loose effect on cohort size and capital stock. Table 5.4 presents some statistics on cohort characteristics for firms entering in good times ($\theta = \theta_g$) relative to firms entering in bad times ($\theta = \theta_b$). For good realizations of the aggregate shock, the number of firms that choose to enter is higher than for bad realizations. On average, the fraction of firms entering in good times is 19% larger than the fraction of firms entering in bad times. On the other hand, each firm that enters tends to do so with a smaller capital stock in good times than in bad times. The capital stock of an entrant in good times is 0.5% smaller than that of an entrant in bad times. Overall, even though each entrant is smaller, the number of firms entering is so much larger in good times that cohorts entering in good times account for a larger fraction of aggregate capital (19.6% larger) and aggregate output (20% larger).

Over time, this pattern does not change. Cohorts entering in good times continue to be smaller (on average) than those entering in bad times (on average they have 5% less capital stock and produce 9% less output). Over their lifetime, they account for a larger share of aggregate output (18% more) and aggregate capital (20% more) than firms entering in bad times. Comparing market failure rates, cohorts entering in good times tend to have longer life spans than those entering in bad times. Lastly, on average, cohorts entering in good times are less prone to market failure than those entering in bad times. These results are reported in Table 5.5.

5.2.3. Impact of Firm Heterogeneity on Aggregate Dynamics

As reported in Doms and Dunne [1998], the distribution of investment is skewed with a small number of firms accounting for a relatively large share of investment. They calculate the Herfindahl index of investment[16] for each year between $1973 - 1988$ and compare this to aggregate investment over the same period. The important finding is that the two series tend to move together with periods of high Herfindahls corresponding to periods of high aggregate investment. This implies that heterogeneity in investment decisions at the firm-level has an impact on aggregate investment in the industry, with periods of high concentration also being periods of high aggregate investment.

The model developed in this paper replicates this feature of the data. The time series for aggregate investment and the Herfindahl index are positively correlated. Periods of high concentration (of investment in the hands of a few firms) are also periods of high aggregate investment in the industry. Thus, heterogeneity among firms does have an impact on aggregate dynamics with the investment decisions of a few firms having a significant impact on industry investment.

5.2.4. Entry, Exit and Aggregate Fluctuations

In this model, periods of high growth in output result in periods of high entry. This matches the fact reported in Hause and Rietz [1984] that the rate of entry is an increasing function of industry growth. Looking at the time series for exit rates, we find that periods of high growth are also periods of high exit. This result appears counter-intuitive as it implies that a larger fraction of firms will exit in periods of high growth. The result does make sense when considered with the fact that new firms have a very high rate of market failure and

[16] The Herfindahl index for investment is the sum across all firms of the squared investment shares, i.e., $\Sigma \left(\frac{I_i}{TI}\right)^2$ where I_i is the investment in firm i and TI is the aggregate investment.

Table 5.5: **Evolution of Cohorts Over Time Conditional on the Aggregate Shock in the Year of Entry**

	1	**2**	**3**	**4**	**5**	**6**	**7**
Market Share	1.19	1.12	1.19	1.21	1.24	1.15	1.22
Average Size of Surviving Firms	0.87	0.92	0.93	0.91	0.99	0.95	0.95
Market Failure Rate	–	0.995	0.833	0.8	0.79	0.79	0.8
Average Capital Stock	0.995	0.97	0.92	0.94	0.92	0.95	0.95
Share of Aggregate Capital	1.2	1.19	1.18	1.22	1.26	1.22	1.19

This table looks at the evolution of cohorts entering during good realizations of the aggregate shock relative to the evolution of cohorts entering during bad realizations of the aggregate shock. Each entry in the table is a ratio reporting the cohort characteristic for cohorts entering in good times relative to those entering in bad times. These values are calculated for seven 5-year periods.

Table 5.6: **Correlation Between Entry Rates and Future Exit Rates: Empirical Evidence versus Implications of the Model**

Empirical Evidence

	Exit Rate in 1972	Exit Rate in 1977	Exit Rate in 1982
Entry Rate in 1967	0.447	0.358	0.237

Results from the Model

	Exit Rate in t	Exit Rate in $t+1$	Exit Rate in $t+2$
Entry Rate in t	0.962	0.532	0.310

This table lays out the correlation between the entry and exit rates in future periods as seen in the data as well as implied by the model. Empirical evidence is drawn from Dunne, Roberts and Samuelson [1988] who calculate these correlations for 387 four-digit industries.

exit within a few periods of entry. As periods of high growth are also periods of high entry, this implies that a large fraction of these new entrants will exit in the subsequent periods. Table 5.6 looks at the correlation between entry rates and future exit rates.

Adjusting the time series on exit for firms that are relatively new to the industry, we find that the correlation between growth in output and the exit rate becomes negative. Thus, for the firms who have survived in the industry for more than two periods, the exit rate is lower in periods of high growth (see Table 5.7). This result is consistent with the finding in Dunne, Roberts and Samuelson [1988] that, after correcting for fixed industry effects, entry rates are negatively correlated with exit rates for any given industry.

Thus, periods of growth (in terms of output) in the industry result in high entry and exit rates, but the exit rates of older firms tends to be negatively correlated with the growth in output.

Table 5.7: **Correlation Between Growth in Output and Exit Rates Conditional on the Age of the Firm**

	Growth in Output
Exit Rate for All Firms	0.3624
Exit Rate for Firms > 1 year old	0.0611
Exit Rate for Firms > 2 years old	−0.1055
Exit Rate for Firms > 3 years old	−0.2025

This table reports the correlation between growth in industry output and the exit rate of firms. It calculates these correlation coefficients for different groups of firms who are differentiated on the basis of age.

Table 5.8: **Correlation Between Growth in Output and the Mean and Variance of Capital Stock and Output Across Firms**

	Growth in Output
Mean Capital Stock	−0.1972
Variance in Capital Stock	−0.1373
Mean Output	−0.2703
Variance in Output	−0.2530

The table above reports the correlation between the growth in industry output and the features of the distribution (of firms across different levels of capital and output) as implied by the model. It calculates the correlation between the mean and the variance of the distribution of capital stock across firms with the growth in industry output. It also calculates the correlation between the mean and the variance of the distribution of output across firms with the growth in industry output.

5.2.5. Industry Evolution: Size and Structure

Due to the entry and exit of firms in equilibrium, industry size fluctuates with the aggregate shock and over the course of the business cycle. Periods of high growth in output are also periods when the industry size is large (in terms of the number of firms in the industry). Studying the change in composition of the industry over time, the pattern that emerges is that in periods of high growth, the size of the average firm tends to be smaller and the firms are spread over a smaller range of capital values, i.e., in periods of growth, the distribution of firms has a lower mean and variance (see Table 5.8).

Putting the facts for industry size and structure together, the implications of this model are that, in periods of boom, the industry size is larger but the average size of firms in the industry is smaller. In periods of recession, the industry size shrinks and the firms that remain in the industry tend to be larger. The correlation coefficient between industry size and average size of a firm in the industry is given in Table 5.9.

6. Conclusion and Extensions

The objective of this paper is to develop and solve a realistic model of industry behavior. Empirical evidence on firm behavior indicates the existence of heterogeneity among firms

Table 5.9: **Correlation Between Industry Size and the Average Capital Stock and Output Per Firm**

	Average Capital Stock	Average Output
Industry Size	-0.9794	-0.7688

This table reports the correlation between the size of the industry and the average capital stock held by a firm in the industry. It also reports the correlation between the size of the industry and the average amount of output produced by a firm in the industry.

and this paper explores the relationship between firm-level heterogeneity and aggregate fluctuations. The model replicates the cyclical patterns of entry and exit observed in the data and provides a framework within which to study the changes in industry size and structure over time. The numerical technique developed in this paper allows us to characterize the equilibrium behavior of the industry under these conditions.

An important application of this model is in the area of policy. It provides a realistic and pertinent framework within which to evaluate the impact of various policy changes on industry dynamics. Changes in fiscal and regulatory policy depend on the distribution of firms and their response to incentives. This model allows us to study the impact of policy changes on the size and structure of the industry and is a useful tool for making policy recommendations.

Extensions to the model include incorporating fixed costs of adjusting capital. As noted by Caballero, Engel and Haltiwanger [1995] and Cooper, Haltiwanger and Power [1995], microeconomic investment decisions tend to be lumpy and discontinuous, implying that the constraints to capital adjustment faced by firms is significantly different from those implicit in the quadratic adjustment costs used in this model. Using non-convex adjustments costs would help replicate the lumpy investment patterns observed in the data. The second extension to the model is to place it in a general equilibrium setting, i.e., in a framework where the demand function is not exogenously determined.

Thus, this chapter looks at the issue of firm-level heterogeneity and evaluates whether heterogeneity at the microeconomic level plays a significant role in aggregate dynamics. Computing the equilibrium for such a model has been a problem due to the lack of an analytical or close-form solution. The use of computational methods allows us to solve and quantitatively analyze a whole range of models which is not otherwise possible.

References

[1] Caballero, R.J., "A Fallacy of Composition", *American Economic Review,* **82** (1992), 1279-1292.

[2] Caballero, R.J., E.M.R.A. Engel and J.C. Haltiwanger, "Plant-Level Adjustment and Aggregate Investment Dynamics", *Brookings Papers on Economic Activity*, (1995), 1-39.

[3] Campbell, J.R and J.D.M. Fisher, "Aggregate Employment Fluctuations with Microeconomic Asymmetries", Working Paper, University of Rochester, NBER and Federal Reserve Bank of Chicago.

[4] Cooper, R., J. Haltiwanger and L. Power, "Machine Replacement and the Business Cycle: Lumps and Bumps", *Manuscript*, (1995).

[5] Doms, M. and T. Dunne, "Capital Adjustment Patterns in Manufacturing Plants", *Journal of Economic Literature*, (1998), 409-429.

[6] Dunne, T., M.J. Roberts, and L. Samuelson, "The Growth and Failure of U.S. Manufacturing Plants", *Quarterly Journal of Economics*, **104** (1989), 671-698.

[7] Dunne, T., M.J. Roberts, and L. Samuelson, "Patterns of Firm Entry and Exit in U.S. Manufacturing Industries", *RAND Journal of Economics*, **19** (1988), 495-515.

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

[9] Evans, D.S., "Tests of Alternative Theories of Firm Growth", *Journal of Political Economy,* **95** (1987), 657-674.

[10] Hopenhayn, H., "Entry, Exit and Firm Dynamics in Long Run Equilibrium", *Econometrica,* **60** (1992), 1127-1150.

[11] Hopenhayn, H and R. Rogerson, "Job Turnover and Policy Evaluation: A General Equilibrium Analysis", *Journal of Political Economy*, **101** (1993), 915-938.

[12] Hause, J.C., and G. Du Rietz, "Entry, Industry Growth and the Microdynamics of Industry Supply", *Journal of Political Economy*, **92** (1984), 733-757.

[13] Krusell, P. and A.A. Smith, Jr., "Income and Wealth Heterogeneity in the Macroeconomy", *Journal of Political Economy*, **106** (1998), 867-896.

[14] Kydland, F.E., and E.C. Prescott, "Time to Build and Aggregate Fluctuations", *Econometrica*, **50** (1982), 1345-1370.

[15] Lucas, R. and E.C. Prescott, "Investment under Uncertainty", *Econometrica*, **39** (1971), 659-681.

[16] Mansfield, E., "Entry, Gibrat's Law, Innovation and the Growth of Firms", *American Economic Review*, **52** (1962), 1023-1051.

[17] Orr, D., "Determinants of Entry: A Study of the Canadian Manufacturing Industries", *Review of Economics and Statistics*, **56** (1974), 58-66.

[18] 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.

[19] Press, W.H., B.P. Flannery, S.A. Teukolsky, and W.T. Vetterling, "Numerical Recipes", Cambridge University Press, Cambridge.

[20] Tauchen, G., "Finite State Markov-Chain Approximations to Univariate and Vector Autoregressions", *Economics Letters*, **20** (1986), 177-181.

[21] Troske, K.R., "The Dynamic Adjustment Process of Firm Entry and Exit in Manufacturing and Finance, Insurance and Real Estate", *Journal of Law and Economics*, **39** (1996), -757.

7. Appendix

This appendix provides a description of the numerical techniques used to solve the firm's dynamic programming problem. The algorithm is similar to the one used in Krusell and Smith [1998]. The difference lies in the fact that the algorithm used in this paper is based on keeping track of not one, but two moments of the capital distribution, namely, K_{t-1}, the total capital stock in the industry in time $t-1$ and N_{t-1}, the number of firms that decide in $t-1$ to enter the industry in time t. The algorithm can be modified to include more moments and for different model specifications.

The objective of the algorithm is to approximate the two functions

$$V(k_{it-1}, z_{it}; N_{t-1}, \theta_g)$$

and

$$V(k_{it-1}, z_{it}; N_{t-1}, \theta_b)$$

We do this by approximating the values of the function on a coarse grid of points in the (k, z, K, N) space.

The firm-specific idiosyncratic shock z follows a continuous-valued autoregressive process of order 1 and can take on any value in $(-\infty, +\infty)$. In determining the grid for the z-shock, we use the technique developed by Tauchen [1986]. This technique uses a discrete-valued Markov chain to approximate the sample path of a continuous-valued autoregressive process. First, a discrete grid for the z-shock is created such that it covers three standard deviations about the mean. Second, the method developed by Tauchen is used to solve for the transition probability matrix of moving from one point on the grid to any other point on the grid, i.e., the probability $F(z', z)$ where both z' and z lie on the discrete grid. The resulting process for z replicates a continuous-valued $AR(1)$ process with the required persistence parameter and unconditional variance.

The grid for k, K and N can be determined in a more flexible way as we do not restrict the choices for capital (individual and aggregate) and fraction of firms entering to points on the grid. Instead, interpolation techniques such as polynomial interpolation and cubic spline interpolation are used to calculate the value function at points not on the grid. Since there is not much curvature in the value function in the K-direction or in the N-direction, we use a small number of grid points. Polynomial interpolation is used to compute the value function for K-values and N-values not on the grid. We compute the value of the interpolating polynomial using Neville's algorithm [17]. In the k-direction, the value function has a point of discontinuity and a fair amount of curvature. The point of discontinuity corresponds to the exit rule described in Section 2 of Chapter 2. Therefore, in this direction, we use cubic spline interpolation which fits a piece-wise cubic function through the given function values, with one piece for each interval defined by the grid. Note that the required cubic splines need to be computed only once for each iteration of the algorithm. Once computed, it is easy to use these splines to calculate interpolated values.

[17] As described in Chapter 3 of Press et al. [1989].

The results produced by this algorithm are not sensitive to increasing the number of grid points in the k, K or N directions. Increasing the number of grid points in the z direction does result in a closer approximation of the continuous-valued Markov process but the gains from doing this are very small.

The numerical procedure can be summed up by the following points:

1. Choose a grid of points in the (k, z, K, N) space using the methods described above.

2. Choose initial values for each of the two functions at each of the grid points.

3. For each of the two values of the θ shock, maximize the right hand side of Bellman equation at each point on the grid. In this maximization, firms are allowed to select any level of capital and various interpolation techniques are use to calculate the value function at points not on the grid.

4. Compare the new optimal values generated by Step 3 to the original values. If the new values are close to the old values, we stop the iteration. Otherwise, Step 3 is repeated until the new and old values are sufficiently close.

5. In order to simulate the behavior of agents, we need to approximate the decision rules associated with the value function computed in Step 4. We approximate the decision rules by first computing optimal decisions on a fine grid of points in the (k, z, K, N) for each value of θ. When computing these optimal decisions, we use the approximate value function as computed in Step 4. Optimal decisions at points not on the grid are then computed using bilinear interpolation (see Press et al. [1989])[18].

[18] For a more detailed explanation of the numerical methods used, see Krusell and Smith [1998].