

The research program of the Center for Economic Studies (CES) produces a wide range of theoretical and empirical economic analyses that serve to improve the statistical programs of the U.S. Bureau of the Census. Many of these analyses take the form of CES research papers. The papers are intended to make the results of CES research available to economists and other interested parties in order to encourage discussion and obtain suggestions for revision before publication. The papers are unofficial and have not undergone the review accorded official Census Bureau publications. The opinions and conclusions expressed in the papers are those of the authors and do not necessarily represent those of the U.S. Bureau of the Census. Republication in whole or part must be cleared with the authors.

THE INDUSTRY LIFE-CYCLE OF THE SIZE DISTRIBUTION OF FIRMS

by

Emin M. Dinlersoz *
University of Houston

and

Glenn MacDonald *
Washington University in St. Louis

CES 05-10 July, 2005

All papers are screened to ensure that they do not disclose confidential information. Persons who wish to obtain a copy of the paper, submit comments about the paper, or obtain general information about the series should contact Sang V. Nguyen, Editor, Discussion Papers, Center for Economic Studies, Washington Plaza II, Room 206, Bureau of the Census, Washington, DC 20233-6300, (301-763-1882) or INTERNET address snguyen@ces.census.gov.

The Industry Life-Cycle of The Size Distribution of Firms*

Emin M. Dinlersoz[†]
University of Houston

Glenn MacDonald[‡]
Washington University in St. Louis

June 2005

Abstract

This paper analyzes the evolution of the distributions of output and employment across firms in U.S. manufacturing industries from 1963 until 1997. The evolutions of the employment and output distributions differ, but display strong inter-industry regularities, including that the nature of the evolution depends whether the industry is experiencing growth, shakeout, maturity, or decline. The observed patterns have implications for theories of industry dynamics and evolution.

JEL Classification: L11, L60.

Keywords: Firm size distribution, industry evolution, industry dynamics, manufacturing industries.

*We thank Roger Sherman and seminar participants at the Universities of Houston and Iowa and Texas A&M University for helpful comments and suggestions. Part of this research was conducted when the first author was a research associate at the California Census Research Data Center (CCRDC) at the University of California at Berkeley. The output in this paper was screened to prevent disclosure of confidential data. The results and views expressed here are those of the authors and do not necessarily indicate concurrence by the Census Bureau. We gratefully acknowledge the assistance of Ritch Milby of the CCRDC. Melanie Fox-Kean provided expert research assistance in the early phases of this project. Financial support was provided by the Center for Research in Economics and Strategy at the John M. Olin School of Business.

[†]Department of Economics, 204 McElhinney Hall, Houston, TX 77204-5019; edinlers@mail.uh.edu

[‡]Olin School of Business, Washington University in St. Louis, Campus Box 1133, One Brookings Drive, St. Louis, MO, 63130-4899; macdonald@wustl.edu

1 Introduction

The size distribution of firms has been the subject of a great deal of theoretical and empirical research; see, e.g., Gibrat (1938), Simon and Bonini (1958), Ijiri and Simon (1964,1974,1977), and Lucas (1978), and more recently, Sutton (1991,1997), Kumar, Rajan, and Zingales (2001), Axtell (2001), and Rossi-Hansberg and Wright (2004). This attention is well-deserved because the size distribution can reveal much about the distribution of productivity, the heterogeneity of production technology, and the degree and type of competition among firms. To industrial organization economists, the importance of understanding changes in the size distribution is akin to the importance of understanding changes in income inequality for growth and development economists, or the importance of understanding changes in wage heterogeneity for labor economists. Describing the evolution of firm size heterogeneity is a critical task for understanding industry evolution and the resulting industry structure.

When all U. S. manufacturing firms are lumped together, the general shape of the size distribution, measured either by employment or value of output, changes little over time; see, e.g., Ijiri and Simon (1964, 1974, 1977), Sutton (1997), and Axtell (2001).¹ This is perhaps surprising since empirical findings on industry life-cycles, theoretical models of industry life-cycles and dynamics, and empirical patterns of firm and industry dynamics, collectively suggest that the size distribution should change as an industry ages.² Evidently, manufacturing is too broad an aggregate for these evolutions to be discernible in the data. Analysis of some narrower definition of an industry is required to determine whether the predicted evolutions are actually there. However, despite the importance of understanding how heterogeneity among firms changes over time, the empirical literature on industry dynamics has paid little attention to the evolution of firm size. In both static and dynamic studies of distribution of firm sizes, with few exceptions, industries are lumped together regardless of whether they are in infancy, in maturity, or in decline. This aggregation may well have obscured important regularities; one of the objectives of our work is to discover whether this is the case.

Gort and Klepper (1982) were the first to establish that manufacturing industries, despite many differences in production technology, inputs, market structures, regulations, etc., go through remarkably similar life-cycle phases. This work was extended and refined by Klepper and Graddy (1990), Agarwal and Gort (1996, 2001), Klepper and Simons (2000), and Simons (2001). A key finding is that with few exceptions, the time path the number of participating firms follows is non-monotonic; a representative path is sketched in Figure 1. An initial rapid rise in the number of firms is followed by a phase called the “shakeout” during which the number of firms falls before eventually becoming stable. Growing output and declining price accompany this non-monotonic path. In addition to the

¹In particular, Axtell (2001) reports a general stability in firm employee size distribution in the form of a Zipf’s Law for the entire population of U.S. firms, not just manufacturing. See also Rossi-Hansberg and Wright (2004) for a study of establishment (but not firm) size dynamics and a discussion on the applicability of Zipf’s Law in their setting.

²For empirical analysis of industry life-cycles, see, e.g., Gort and Klepper (1982) and Agarwal (1998). For theoretical models of industry evolution, see, e.g., Jovanovic (1982), Jovanovic and MacDonald (1994a,b). For empirical patterns of firm growth, entry and exit, see, e.g., Dunne, Roberts and Samuelson (1988, 1990).

phases in Figure 1, a decline/contraction life-cycle phase is increasingly common in manufacturing; in this phase both the number of firms and output decrease. A second key conclusion of these papers is that life-cycle movements in price, output, and the number of firms are generally much stronger than business cycle effects or other industry-wide economic shocks, and determine the long-run trends in an industry.

Our objective in this paper is to describe, in as much detail as the data allow, how firm heterogeneity evolves as an industry goes through its life-cycle. For this purpose, we use the most comprehensive U.S. data on firm sizes: the Census of Manufactures. These data include a large number of manufacturing industries – we are able to study 322 4-digit industries – and cover a 35 year period including 1963, and 1967 through 1997 at five year intervals. We document industry life-cycle changes in the moments of firm size, i.e., mean and median, standard deviation, coefficient of variation, skewness and kurtosis, as well as general shifts in the firm size distribution. By examining the behavior of the entire firm size distribution and its moments, as opposed to a summary measure such as a concentration ratio, we offer more comprehensive evidence on the evolution of firm heterogeneity. Moreover, our focus on higher moments of the distribution is driven by both industry life-cycle theory and other distribution-oriented work (e.g., on income distribution, or the returns to firm investments) suggesting that much of the impact of firm choices reveals itself in the tails of the distributions of outcomes such as firm size or profitability.

Our main findings, which we demonstrate in several ways, are that the firm size distribution, with size measured either by employment or output (not sales), exhibits (*i*) remarkable inter-industry regularity, (*ii*) significantly different behavior depending on which life-cycle phase the industry is experiencing, i.e., growth, shakeout, stability, or decline; and (*iii*) significantly different behavior depending on whether size is measured by output or employment. We also find support for the idea that the evolution of the size distribution is more dramatic in industries where firm growth evolves more dramatically. Finally, we assess the consistency of existing models of industry evolution with our empirical findings. The models' general implications that the size distribution should display increasing mean or median output, and that asymmetry in the size distribution should evolve non-monotonically – in particular, mirroring growth in firm numbers – are both quite evident in the data. At the same time, the models' other important implication, that heterogeneity will tend to disappear after firm numbers peak, fares less well. Indeed, in these data even the coefficient of variation rises throughout the life cycle, albeit more quickly during the earlier phases of the life cycle.

An important caveat to all these conclusions is that given the brevity of the period covered by our data, most industries are dominated by one life-cycle phase. Thus, when comparing industries with growing firm numbers to those with shrinking numbers, we are also comparing two different sets of industries. And so, despite the encouraging Gort-Klepper result that most industries display common evolutionary patterns, we cannot be sure that when one of our growing industries begins to shrink, it will shrink in a manner similar to the shrinking industries in our data.

The general patterns we find are as follows. Consider output first. The distribution of output is

highly positively skewed throughout the life-cycle. Mean and median output increase as the industry ages, with the increase most rapid during the shakeout, consistent with the familiar notion that exit is dominated by smaller firms who have somehow failed to do what the growing firms have done. The standard deviation also rises through all phases, but much more during the phase dominated by entry, consistent with entry producing firm heterogeneity whose growth is lessened during the shakeout. The growth in heterogeneity dominates the growth in firm size in the sense that the coefficient of variation rises through the life-cycle. The evolution of higher moments varies more by industry phase. During the rapid entry phase, firm output becomes much more positively skewed and “fat-tailed” (i.e., increasing kurtosis) – altogether a more asymmetric distribution. This evolution also appears, but much less dramatically, during industry maturity, i.e., when firm numbers stabilize but total industry output continues to grow. Thus, periods of either growth or stability in firm numbers display increasing asymmetry in output, with the degree depending on the extent of growth in firm numbers. In contrast, when the industry goes through the shakeout – falling firm numbers alongside growing industry output – the growth in skewness and kurtosis is reversed, yielding a less asymmetric firm size distribution. But the decline in asymmetry is less dramatic than the increase that occurred when firm numbers were growing. In fact, the tendency towards less asymmetry is more pronounced during the period of industry decline, i.e., when both firm numbers and output are shrinking, than during the shakeout. Evidently, the forces working to produce greater asymmetry are always at work when the industry is growing.

When size is measured by employment, the evolution of asymmetry is similar to what is observed when size is measured by output. But whereas mean or median output is always growing, mean or median employment shrinks rapidly during the entry phase and grows even more rapidly during the shakeout. This is consistent with both observation and theory suggesting both entering and exiting firms are smaller than incumbent or continuing firms. There is little evolution of mean or median employment during either maturity or decline. A similar pattern emerges in terms of standard deviation, i.e., falling during the rapid entry phase and rising during the shakeout. Interestingly, the movements in heterogeneity and mean employment always lead to greater coefficient of variation. That is, when mean employment grows, heterogeneity grows faster. And when mean employment shrinks, heterogeneity shrinks less.

We also explore changes in the size distribution in the sense of first order stochastic dominance. Interestingly, most of the empirical size distributions can (in a sense we make precise below) be described as either increasing or decreasing. When size is measured by output, all but a small fraction of industries experiencing the shakeout also display an increasing size distribution. But in all other phases, increasing size distributions are more frequent than decreasing, but much less so than in the shakeout phase; for example, 44.5% versus 37% in the rapid entry phase. When size is measured by employment, roughly half of industries experiencing the shakeout also display an increasing distribution. In all other phases, a decreasing size distribution is much more frequent; for example, 62.9% versus 19.6% in the rapid entry phase.

Finally, we show that the more severe the life-cycle effects, as measured by the magnitude of the change in firm numbers, the greater the growth in the moments of the size distribution, however size is measured. More dramatic evolution of firm numbers is associated with more dramatic evolution in the size distribution.

Our paper is related to a small but growing empirical literature on the dynamics of the firm size distribution. Cabral and Mata (2002) analyze the evolution of the firm size distribution by following cohorts of entrants in Portuguese manufacturing industries over a period of eight years. The logarithm of firm size (employment) for entering firms in these industries is positively skewed initially and becomes more symmetric as a result of firm growth and exit. Lotti and Santarelli (2004) also examine the evolution of firm size in entering cohorts in five Italian manufacturing industries over a period of six years. The logarithm of employment in these industries is also initially positively skewed and tends to converge to a limit distribution over time. We improve on this body of work by studying a greater number of industries over a longer time period in a much larger economy. Also, consistent with what is suggested by both theory and data on industry life-cycles, we emphasize the different evolutions that occur during the various life-cycle phases. Finally, by using both employment and output as measures of size, we are able to uncover more about the processes leading to output growth. For example, significant mean output growth can be sustained alongside significant mean employment decline.

The rest of the paper is organized as follows. Section 2 lays out the theoretical motivation. The data and empirical methodology are described in Sections 3 and 4. Selected industries are studied in Section 5 to illustrate the empirical approach. This is followed by the full analysis in Section 6, including some robustness checks. Section 7 concludes.

2 Theories

2.1 Models of industry evolution

The evolutionary trends displayed in Figure 1 were first documented by Gort and Klepper (1982), and later extended by Klepper and Graddy (1990), Agarwal (1998), Gort and Agarwal (2001), and others. Several models of industry evolution have been employed to explain these trends. Jovanovic and MacDonald (1994a) offer predictions about the evolution of the firm size distribution. The basic setup is as follows. The industry is born when a technological innovation allows for “low-tech” production, at which time some firms enter and begin production. Later, a major technological “refinement” arrives. Existing low-tech firms can try to adopt it and become “high-tech.” New firms can also enter with the hope of innovating. Some of these newcomers innovate and become low-tech (it is assumed that low-tech know how is required to implement high-tech). Of this low-tech group, those that later fail to become high-tech may exit. That is, once the refinement has arrived, in each period a constant fraction of the existing low-tech firms innovate and become high-tech. Thus, gradually, an increasing fraction of firms becomes high-tech. Technology dictates firm size, and high-tech firms are assumed to

have lower marginal cost, and hence higher output. As a result, until exit begins, industry output must rise and price must fall as the mixture of firms shifts towards high-tech. In the rest of the industry's life-cycle two things can happen. In one scenario, price does not fall enough for low-tech firms ever to choose to exit. In the second scenario, product price falls enough for exit to occur. If the high-tech firms are much larger than low-tech firms or if it is easy to become a high-tech firm, price falls and output rises very quickly, implying a mass exit of low-tech firms. Otherwise, exit is gradual. Gradual and mass exit are both empirically relevant, as is no exit.³

This simple model generates testable implications about the time path of the distribution of output. At the outset, the distribution is degenerate at the output of a low-tech firm. As the mix of the firms shifts towards high-tech, price falls, depressing the sizes of all firms. But at the same time, the increasing fraction of high-tech firms acts to increase average firm size. Depending on the strength of the two effects, the firm size distribution can initially stochastically decrease or increase.⁴ Once the exit of low-tech firms starts, price stabilizes and firm size stochastically increases as the mix of firms shifts further towards high-tech. Output becomes more dispersed initially as some firms become high-tech, but the dispersion eventually declines as the fraction of high-tech firms increases. Firm size is also positively skewed when there are only a few adopters of high-tech know-how, but skewness declines as the industry matures. Similarly, kurtosis is also high early on when a small fraction of firms are high-tech, but declines as the mix of firms changes in favor of high-tech firms, and eventually increases again when the industry is made up of mostly high-tech firms. These implications are derived in Appendix A. Using the Gort and Klepper data for the U.S. automobile tire industry, the expressions in the Appendix, and the estimated parameters from Jovanovic and MacDonald (1994a), the evolution of the moments of firm size distribution can also be estimated; they are displayed in Figure 2.⁵

Jovanovic and MacDonald (1994b) consider a richer setup for firm heterogeneity. While this model abstracts from entry and exit, it provides further insight into the evolution of the firm size distribution. In that model, greater technological know-how improves a firm's technology and lowers its total and marginal costs. Firms are endowed with some know-how, but can also improve it through innovation and imitation, both of which are costly and imperfectly controllable. Differential innovative and imitative success causes the distribution of know-how, and hence firm outputs, to continue to evolve

³See Gort and Klepper (1982) and Agarwal (1998).

⁴Throughout the paper, "stochastically increasing (decreasing)" refers to an "increase (decrease)" in the sense of first-order stochastic dominance.

⁵The model described above does not fully incorporate heterogeneity across firms. The fact that there are only two types of firms simplifies, but also gives a special structure to the model. In addition, the episode of entry by new firms when the technological refinement becomes available is represented by a single period, which may actually correspond to several years of entry in data. If entry continues as some firms become high-tech, the firm size distribution can decline stochastically, especially if the entry rate is high and the rate of adoption of high-tech know-how is low. Another simplification in the model is the instantaneous firm growth. Most firms are actually small upon entry and grow slowly, and may not achieve their optimal size for a while. This gradual growth of firms may reinforce the tendency of firm size to decline stochastically as entry occurs. Similarly, it can also slow down the stochastic increase in firm size as adopters of high-tech production methods grow only gradually.

over time. Average firm size increases, but firm size does not necessarily increase stochastically. The variance of firm size generally exhibits a non-monotonic path, initially increasing as firms become more diverse in their know-how due to innovation, and less diverse as firms take advantage of the fact that heterogeneity in know-how opens up opportunities for imitation. Eventually, if there is an upper bound on technological know-how, firms may become technologically similar as opportunities to innovate further dwindle, and imitation by firms whose know-how is less advanced continues. However, depending on the relative costs of innovation and imitation, as some firms expand the frontier of know-how and others try to catch-up, the non-monotonic path of variance may repeat itself before the distribution of know-how settles, if ever. A similar argument applies to skewness and kurtosis.

Klepper (1996) also features industry evolution based on technological innovation. In his model, firms differ in their success in innovation, and can engage in both product and process innovation, the former aimed at introducing a new product, and the latter reducing the cost of production. In any period, entry and exit can occur, and all firms engage in one or both types of innovation. Over time, incumbents grow and it becomes harder for new entrants to surmount incumbents' scale advantages through new product introductions. Gradually, process innovation starts to overtake product innovation, and entry eventually ceases. Less innovative entrants exit as others continue to expand. Over time, firm size can stochastically increase or decrease, depending on the strength of the entry and expansion effects. As exit takes over and the number of firms declines, the expansion of the firms remaining in the industry, coupled with exit by smaller firms, leads to a stochastic increase in firm size and to an increase in average firm size. The variance of firm size can increase in the entry phase as small firms enter and coexist with expanding incumbents. However, as exit occurs and smaller firms exit, firm size becomes less dispersed and variance declines. Skewness is positive early on when smaller firms enter and existing firms expand, but can become negative later as smaller firms exit and the remaining firms become larger. Kurtosis can also exhibit a non-monotonic path, increasing initially as entry leads to a more peaked distribution and firm expansion thickens the right tail of the distribution, but decreasing later as smaller firms exit.

Some of the implications of the models discussed so far can also emerge in other models that abstract from technological innovations. For instance, in Jovanovic's (1982) model of selection and industry evolution, the size distribution evolves as firms gradually learn about their intrinsic efficiency, and those that discover they are inefficient exit. Average firm size can increase or decrease, depending on the behavior of prices. Initially, for a given cohort of firms, all entrants are the same size, so the variance of firm size is zero. But, over time, heterogeneity increases as firms' outputs diverge; however, this increase need not be monotonic. Ericson and Pakes (1995) also consider a rich model of industry dynamics which can accommodate a variety of models of competition. The implications of their model for the evolution of the moments of the size distribution depend on the exact specification of the type of competition among firms and no general predictions can be made. The implications of the various models discussed so far are summarized in Table 1.

Existing empirical regularities about firm growth and turnover also have implications for the evo-

lution of firm size distribution. Two observations are especially relevant. First, entrants are usually small and grow slowly.⁶ Second, smaller and younger firms are more likely to exit compared to older and larger firms.⁷ Thus, average firm size is expected to decrease during the rapid entry phase of the industry life-cycle, and then increase during the shakeout phase as exit dominates entry.

2.2 Declining industries

The models discussed thus far assume enough demand to allow operating firms to continue to grow, or at least not to shrink persistently. But, obviously, an industry can experience a decline or contraction phase if demand shrinks persistently due to, e.g., the availability of a superior substitute or the obsolescence of the product. Chemical industries offer numerous examples of this type of decline; see, e.g., Ghemawat and Nalebuff (1990) and Lieberman (1990).

Dynamic models of competitive industries, such as Hopenhayn (1992), predict that as demand shrinks and price declines, the least efficient producers are the first to exit. As long as firm size and efficiency are positively correlated, the firm size distribution must then stochastically increase as smaller firms exit and the left tail of the distribution is trimmed. When strategic interaction between firms is important, however, larger firms can exit first, as shown by Ghemawat and Nalebuff (1985, 1990) and Whinston (1988). Overall, theories of industrial decline offer a variety of predictions about the evolution of firm size distribution.

Sutton (1997) recently assessed the literature on industry decline in light of evidence on the evolution of 4- and 8-firm concentration ratios for the industries that exhibited a net loss in the number of firms exceeding 40%. There appears to be little systematic change in concentration ratios in this sample. This investigation can be improved in two dimensions. First, Sutton's (1997) analysis probably confuses industries that went through their shakeout phase with those that genuinely declined; for example, in Gort and Klepper's data, the average industry that reached maturity experienced a roughly 40% decline in firm numbers during the shakeout. Since the economic forces driving industry evolution are likely to be quite different during the shakeout and decline phases, identification and separate analysis of the two phases should be useful and may help in explaining Sutton's findings. Below, we distinguish between shakeout and decline by noting that whereas the number of firms exhibits a similar pattern in both phases, industry output increases during shakeout but decreases during decline. Second, by investigating the behavior of the entire size distribution, rather than just a summary measure such as the 4- or 8-firm concentration ratio, we expect to deliver a more complete picture of the behavior of the firm size distribution.

⁶See Geroski (1995) for a discussion of evidence on the general properties of entrants. See Dunne, Roberts, and Samuelson (1988) for evidence on the size of entrants in U.S. manufacturing industries.

⁷See Dunne, Roberts, and Samuelson (1988, 1990).

3 Data

The primary dataset we explore is the U.S. Census Bureau’s Census of Manufactures for 1963, and 1967 through 1997 at five year intervals. More detail on this dataset is provided in Appendix B.1. Two important issues are the definition of an industry and the measurement of firm size.

3.1 Industry definition

Models of industry evolution usually focus on a homogenous product or a group of products that are very closely related. The 1987 industry classification system (SIC), which we employ consistently, consists of 5 levels of aggregation for individual products.⁸ A “product” is defined uniquely by a 7-digit code. Similar products are grouped into “product classes” identified by their common 5-digit code; there were 1,446 such classes in the 1987 SIC system. These product classes are further grouped into 459 “industries” according to the first 4 digits of the product class code.⁹

We focus on the 4-digit level of aggregation. There is no perfect level of aggregation. Our decision to adopt a 4-digit definition reflects an attempt to balance two forces. Too broad a definition, e.g. 2-digit, will obscure the industry life-cycle effects about which the theory makes predictions. Too narrow a definition, e.g. 7-digit, will cause industry evolution to coincide with the evolution of a particular product, e.g., manganese dioxide batteries versus silver oxide batteries, which are arguably very closely related products but nevertheless classified as two separate 7-digit codes. The 4-digit definition appears, based on the analysis below, to balance these forces adequately.¹⁰

Our sample consists of 322 out of a total of 459 4-digit industries defined by the 1987 SIC system. Dropping some industries appears to be necessary to maintain data quality. There are three issues. First, other researchers – Dunne, Roberts and Samuelson (1988) discuss this at length – have uncovered quantitatively significant problems in some product codes for the census years 1963 and 1967. Next, some industries are simply a heterogenous collection of products that are not elsewhere classified. Finally, some industries had missing observations, and others exhibited substantial discontinuities in the time series for number of firms and output that appear to be due to revisions in SIC codes. Appendix B.1 provides the specifics of the selection process that led to our sample of 322.

⁸As discussed in the Appendix, we used the data already constructed by the Census Bureau that is based on a matching of industries over time using the 1987 SIC product classification system.

⁹On average, there were 3.15 5-digit product classes within a 4-digit industry in the 1987 SIC system. This average was highest (5.00) for the 4-digit industries classified under the 2-digit group Printing and Publishing industries, and lowest (1.09) for the 4-digit industries classified under the 2-digit group Leather and Leather Products. A full list of 7-digit and 5-digit product groups classified under each 4-digit industry is available from the U.S. Census Bureau’s 1987 supplement publications to the 1987 Census of Manufactures.

¹⁰The 4-digit level industry classification is generally coarser than the product level analysis of Gort and Klepper (1982). However, even a narrowly defined product category can consist of several products. For instance, the fluorescent lamp, one of the products used by Gort and Klepper (1982), is essentially a product category and contains many different types of lamps in various sizes, shapes, and capacities.

3.2 Plants, firms, and measures of firm size

Theories of industry life-cycles do not distinguish between firms and plants, and treat the industry as a primitive defined by a homogeneous product. However, in the data, most firms have several plants and operate in several industries, depending on how an industry is defined. To deal with plants versus firms, we aggregate plant-level data to firm-level using firm identifiers assigned to each plant. The precise manner in which we do this interacts with our treatment of industries; i.e., we first classify plants into industries, then aggregate to the firm level.

We assign plants to industries in two ways. The first is the “primary-SIC-code-based classification,” which assigns a plant to a 4-digit industry if the plant’s highest value of shipments among all products it produces falls into that industry. This approach is the main reporting method the Census Bureau employs when classifying plants into 4-digit industries, effectively assuming that each plant is single-product rather than a multi-product. This may either understate or overstate the number of firms that, from the standpoint of the theory, should be included in an industry. Following this method, all of a plant’s employment and value of shipments (which we later employ to calculate output) is assumed to belong to its primary SIC code. Our second approach takes into account all the 4-digit industries in which a plant operates, and so considers each plant to be a multi-product. The classification of a plant into each 4-digit industry is done using product level data that provides the value of shipments for each plant by 7-digit product category, which we aggregate to the 4-digit level and then allocate employment and (trivially) value of shipments using these industry shares. Thus, for example, if the value of a plant’s shipments in a 4-digit industry is $x\%$, we assume $x\%$ of the plant’s employment belongs to that industry. We report our results for both classification schemes.

The ideal theoretical measure of firm size is output, rather than employment or sales. While both employment and sales have traditionally been used as measures of firm size, relatively little is known about the relationship among different measures. If firm productivity improves over the industry life-cycle, and especially if the increase is non-uniform across firms of different sizes, firm employment is likely to be a poor measure of heterogeneity in firm size heterogeneity. Sales, on the other hand, suffer from the effects of price changes over time. That is, while output usually has a predictable growth path throughout the industry life-cycle, sales may increase or decrease over time depending on the price elasticity of product demand. The analysis of sales data is further complicated by inflation. To avoid these issues, we construct a measure of output as follows. Firm output is obtained from the total value of shipments of a firm’s plants using 4-digit industry price deflators; these deflators are available from the NBER/CES Manufacturing Productivity Database, described in Appendix B.2. The price deflator for an industry allows us to construct a time series for firm value of shipments in 1987 dollars. Thus, under the assumption that the industry price is common to all firms in a given census year, we can identify the output of a firm in a given 4-digit industry up to a constant. The details of the construction of firm output and industry price deflators are discussed in Appendix B.1.2. The NBER/CES dataset was also used to obtain industry real price series, as described in Appendix

4 Empirical methodology

Our empirical analysis comprises three main steps. First, industries are classified according to the life-cycle phase(s) they display during the sample period. Second, we analyze changes in the firm size distribution during these life-cycle phases. Finally, we document differences in the behavior of the size distribution during different life-cycle stages.

4.1 Identification of life-cycle phases

Gort and Klepper (1982) originally identified the 5 life-cycle “stages” shown in the upper panel of Figure 1. The bulk of the literature building on Gort and Klepper’s work focuses on the three “phases” shown in the lower panel of Figure 1. We do the same. There are two main reasons for our focus on three phases rather than five stages.¹¹ First, in our data, firms are observed primarily at 5-year intervals. This makes a fine-tuned identification of stages impractical. Second, Gort and Klepper’s (1982) stage I had vanished by the early 1960s, roughly when our sample period begins.¹² Third, given the low frequency with which firms are observed, it is inherently difficult to identify the transitory period of stability in firm number preceding the shakeout (stage III). Thus, our phase I is the initial growth phase during which the number of firms in the industry increases, corresponding to stages I and II, and early parts of stage III; phase II is the shakeout during which the number of firms decreases, corresponding to the latter part of stage III, all of stage IV, and early parts of stage V; and finally, phase III is the phase of stability or maturity, corresponding to stage V, during which the number of firms does not change substantially. When an industry is observed to decline, we treat this as a separate terminal phase.

It is important to note that the life-cycle stages or phases sketched in Figure 1, while typical, need not occur in every single industry. For instance, there are some industries that never experienced a shakeout.¹³ So our phase III also includes the case where an industry does not experience a shakeout, but only stability in the number of firms following phase I. Our approach does not rely on all three phases being observed.

Given what is already known about life-cycles, the available time series will rarely be long enough to observe the entire life-cycle of an industry. Instead, we will observe a 35-year-long subset of the life-cycle. In any case, we need to identify the trends in the number of firms and output during these 35 years in order to classify the observed episode into life-cycle phase(s). There are many ways this

¹¹Gort and Klepper (1982) explain that the number of stages they identify is not definitive, and can depend on the nature and the frequency of the data, as well as a researcher’s goal.

¹²See Agarwal and Gort (1998) for evidence on the gradual disappearance of this phase.

¹³Gort and Klepper (1982) found that there was little or no shakeout in the baseboard radiant heater, electrocardiograph and fluorescent lamp industries.

might be done. One simple approach that seems to work well in our data is to identify the underlying trend in the number of firms using a time-series filter. Denote the number of firms in the industry at time t by N_t , for $t = 1, 2, \dots, T$. Treating N_t as a continuous variable for the time being, it is assumed to follow the process

$$N_t = N_t^* + \varepsilon_t,$$

where N_t^* is an underlying “smooth trend” function of t that describes the life-cycle behavior of the number of firms, and ε_t is a zero-mean error component that captures deviations from this trend. Following Hodrick and Prescott (HP, 1997), the trend is the solution to the optimization problem

$$\min_{\{N_t^*\}_{t=1}^T} \left\{ \sum_{t=1}^T \varepsilon_t^2 + \lambda \sum_{t=2}^T [(N_{t+1}^* - N_t^*) - (N_t^* - N_{t-1}^*)]^2 \right\},$$

where $\lambda > 0$ is a parameter that penalizes variability in N_t^* .¹⁴ We use the same procedure to uncover the trend, Q_t^* , in industry output, Q_t . With the estimated trends \widehat{N}_t^* and \widehat{Q}_t^* in hand, the life-cycle phases can be identified from the joint behavior of \widehat{N}_t^* and \widehat{Q}_t^* . If \widehat{Q}_t^* is increasing and \widehat{N}_t^* is increasing (decreasing), then the industry is in phase I (phase II). If \widehat{Q}_t^* is increasing and \widehat{N}_t^* is relatively stable or exhibits no clear trend, then the industry is in phase III. If both \widehat{N}_t^* and \widehat{Q}_t^* are decreasing, then the industry is in its decline.

Figure 4 contains several sample paths for the number of firms. Restrictions imposed by the Census Bureau prohibit us from presenting much detailed industry data. Thus, the sample paths in the figure are artificial but representative of what we observed in most of our sample of industries. In most cases, the number of firms did not fluctuate greatly around the trend, so the trend was easily identified. In others, the number of firms fluctuated more, possibly due to business-cycle effects. When this occurred, we relied more on the HP-filter to determine the underlying trend. In the few cases in which classification was not obvious, we used several different values for the smoothing parameter (λ) to make sure that the classification was made as accurately as possible. While no classification method is guaranteed to be error-free, in most cases the trends were obvious and strong. Overall, we found that for most industries, the trend in number of firms for the entire period of 35 years could be classified as either phase I, phase II, phase III, or phase I combined with phase II. More detail on the classification is provided below.

Samples of time-paths for output are shown in Figure 5. These examples were produced using the value of shipments and price deflator data from the NBER/CES productivity database, which is publicly available. Output data computed from the Census of Manufactures, which we use for our analysis, exhibits similar behavior, but is only available in census years and are subject to disclosure restrictions, especially for industries with small number of firms and high concentration ratios. Since

¹⁴A practical issue is the choice of the smoothing parameter λ . Arguments in Ravn and Uhlig (2002) can be used to obtain a range of λ values for quinquennial data we use. For each industry we experimented with several values of λ in the range $[1, 5]$ and found that $\lambda = 2.5$ to 3 usually worked well in most cases.

the NBER/CES data have a higher frequency (annual) and a longer time span, we chose to use it to generate figures. But to avoid any discrepancies between the two data sets we did not use it to construct the output figures actually used in our empirical analysis.¹⁵ As occurred when we looked for trends in firms numbers, the output generally exhibited three distinct trends throughout the sample period of 35 years: increasing, decreasing, and relatively stable. In a few cases, there was also a non-monotonic (inverted U-shaped) pattern; we also discuss this further below.

4.2 Analysis of the size distribution

Given our classification of industries into phases, we performed a series of statistical analyses the firm size distribution.

4.2.1 Moments of firm size

Let X_t be a random variable that represents firm size in an industry at time t , and let $F_t(x)$ be its distribution function. Throughout the rest of the paper, we use the following definitions: mean $\mu_t \equiv E[X_t]$; median $m_t \equiv \inf\{x : F_t(x) \geq 0.5\}$; standard deviation $\sigma_t \equiv (E[(X_t - \mu_t)^2])^{1/2}$; coefficient of variation $cv_t \equiv \frac{\sigma_t}{\mu_t}$; skewness $\gamma_t \equiv \frac{E[(X_t - \mu_t)^3]}{\sigma_t^3}$; and kurtosis $\kappa_t \equiv \frac{E[(X_t - \mu_t)^4]}{\sigma_t^4}$. Estimators of these moments are:

$$\begin{aligned} \hat{\mu}_t &= \frac{1}{N_t} \sum_{j=1}^{N_t} x_{jt}, \quad \hat{m}_t = \{x_{(N_t/2+1)} \text{ if } N_t \text{ is odd, } \frac{x_{(N_t/2)} + x_{(N_t/2+1)}}{2} \text{ otherwise}\}, \quad (1) \\ \hat{\sigma}_t &= \left(\frac{1}{N_t - 1} \sum_{j=1}^{N_t} (x_{jt} - \hat{\mu}_t)^2 \right)^{1/2}, \quad \hat{cv}_t = \frac{\hat{\sigma}_t}{\hat{\mu}_t}, \\ \hat{\gamma}_t &= \frac{1}{N_t} \sum_{j=1}^{N_t} \left(\frac{x_{jt} - \hat{\mu}_t}{\hat{\sigma}_t} \right)^3, \quad \hat{\kappa}_t = \frac{1}{N_t} \sum_{j=1}^{N_t} \left(\frac{x_{jt} - \hat{\mu}_t}{\hat{\sigma}_t} \right)^4. \end{aligned}$$

The coefficient of variation is important for our purposes because, in many industries, the mean and the standard deviation both evolve in the same direction. When this happens it is informative to have a measure of whether average size is evolving more or less quickly than heterogeneity. Skewness captures the degree of asymmetry in the firm size distribution. A positively skewed distribution tends to have a longer right tail and a mean that exceeds the median; i.e., it has a smaller proportion of the sample that is above them mean, but these observations are especially large. Kurtosis measures the thickness of the tails of the distribution. A higher value of kurtosis means a less peaked, flatter distribution with more of the sample in the tails. Our focus on higher moments of the distribution of

¹⁵We compared the output and employment figures at the 4-digit industry level for the Census of Manufacturers and the NBER/CES database. We found that these measures did not match perfectly across the two datasets. The reason is that the NBER/CES database does not use the census data at the firm level directly to calculate these measures. See Appendix B and Bartelsman and Gray (1996) for details.

firm size is not statistical/mechanical. Instead, it reflects the fact that the theory discussed above, as well as many other empirical studies of differential economic success, e.g., income distribution, suggest that it is likely that we will observe skewed distributions with “fat tails”. Thus, our methodology should allow for these features to be observed readily if they are there.

4.2.2 Stochastic trends

In addition to the changes in individual moments of firm size, an important question is whether the size distribution as a whole is changing significantly over time. If it is, one can go further and ask whether, for example, the change is of the first order stochastic dominance variety.

To look for a general change in the distribution, we proceed as follows. For dates $t = 1$ and $t = T$, we test the hypotheses

$$H_o : F_T(x) = F_1(x) \text{ for all } x. \quad (2)$$

$$H_a : F_T(x) \neq F_1(x) \text{ for some } x.$$

We use relatively flexible methods capable of capturing movements in the size distribution regardless of the exact type of change the distribution is undergoing. A distance-based measure that can be used to test this hypothesis is the Kolmogorov-Smirnov (*KS*) test.¹⁶ Define the empirical counterpart of F_t at any point x in the support of the firm size distribution by

$$\widehat{F}_t(x) = \frac{1}{N_t} \sum_{j=1}^{N_t} I(x_{jt} \leq x),$$

where $I(\cdot)$ is the indicator function. Let S_t be the set of observed firm sizes at time $t = 1, T$. The *KS* test-statistic is given by

$$D = \max_{x \in (S_1 \cup S_T)} \left| \widehat{F}_T(x) - \widehat{F}_1(x) \right|. \quad (3)$$

An attractive feature of the statistic D is that its distribution does not depend on the underlying true distribution of firm size.¹⁷ This property is particularly useful for our purposes because the shape of the size distribution varies from one industry to another, as well as over time. An assumption underlying the *KS* test is that the two samples are independent. This assumption can be violated in

¹⁶A chi-square test is also feasible. However, the *KS* test has certain advantages over the chi-square test. First, the *KS* test does not require data that come in groups or bins, while the performance of the chi-square test is affected by the number of bins and their widths. Second, the *KS* test can be applied for small sample sizes, whereas the chi-square test is more appropriate for larger samples. For more on the *KS* test, see, e.g., Gibbons (1971) and Siegel and Castellan (1988).

¹⁷The critical values of the *KS* test statistic are available in standard texts on nonparametric statistical analysis as well as in common statistical software. See, for instance, Tables L_I to L_{III} in Siegel and Castellan (1988), pp. 348 to 352. We used STATA to calculate the *KS* statistics values and their significance. STATA allows for a better approximation of critical values of the *KS* statistic for small samples. We used this improved approximation when the number of firms at any time of measurement ($t = 1$ or T) was less than 100.

our setting, since the set of firms active at time $t = T$ is likely to contain firms that were in business at $t = 1$. The fact that we measure the size distributions at two points in time that are far apart (35 years) alleviates the concerns about dependence to some extent, but certainly does not eliminate it.¹⁸

The second issue of interest is the direction of change in the size distribution, in a first-order stochastic sense.¹⁹ A one-sided version of the *KS* test can be used to test for first-order stochastic dominance. Define

$$D^- = \max_{x \in (S_1 \cup S_T)} (\widehat{F}_T(x) - \widehat{F}_1(x)) \quad (4)$$

for testing H_o in (2) against the alternative

$$H_a : F_T(x) \geq F_1(x), \text{ for all } x, \text{ and } F_T(x) > F_1(x) \text{ for some } x.$$

Similarly define

$$D^+ = \max_{x \in (S_1 \cup S_T)} (\widehat{F}_1(x) - \widehat{F}_T(x)) \quad (5)$$

to test against the alternative

$$H_a : F_T(x) \leq F_1(x), \text{ for all } x, \text{ and } F_T(x) < F_1(x), \text{ for some } x.$$

A sufficiently large positive value of D^- is consistent with a stochastically decreasing firm size between $t = 1$ and $t = T$. On the other hand, a sufficiently large positive value for D^+ favors a stochastically increasing firm size. However, note that D^- and D^+ can be large simultaneously. This can happen, for instance, if the two distribution functions cross at a single point and the maximum distances between them on both sides of this point are large. To identify cases in which only one of D^- and D^+ is significantly large, we use the following approach: F_T stochastically dominates F_1 if D^+ is statistically significant at $\alpha\%$ or lower, and at the same time, D^- is not significant at $\alpha\%$ or lower. The analogous definition is employed where F_1 stochastically dominates F_T .

4.3 Life-cycle effects

After identifying the life-cycle phases and obtaining the statistics pertaining to the size distribution, we summarize the behavior of the size distribution by life-cycle phase. We proceed as follows. For each industry, let Δy denote the percent growth rate for the empirical moment $-y = \widehat{\mu}, \widehat{m}, \widehat{\sigma}, \widehat{c\hat{v}}, \widehat{\gamma}, \widehat{\kappa}$ – between the two end points of our sample, 1963, corresponding to $t = 1$, and 1997, corresponding

¹⁸Methods have been recently developed to obtain consistent *KS* test statistics under general dependence of the two samples (see, e.g., Linton, Maasoumi, and Whang (2003)). However, they are computationally demanding, so we did not implement them for this analysis.

¹⁹Recently-developed techniques allow higher orders of dominance to be explored. However, the theories discussed earlier do not appear to have obvious predictions on these higher order shifts, so we focus only on first order dominance.

to $t = T$. Denote the expected value of Δy by $\mu_{\Delta y}$, which can be viewed as the mean of an underlying random process that generates the growth rates Δy . For each life-cycle phase, we compute the inter-industry average behavior of Δy , and test the hypotheses

$$H_o : \mu_{\Delta y} = 0,$$

$$H_a : \mu_{\Delta y} \neq 0.$$

We also compare the average value of Δy across phases, and identify any asymmetries in the behavior of the moments across phases. Similarly, we summarize the patterns of stochastic movements in firm size by life-cycle phase based on the *KS* tests.

Another issue of interest is whether the magnitude of the change in each moment of the size distribution is systematically related to the extent of the life-cycle phase; e.g., do the moments move more quickly in industries where the number of firms is evolving more quickly? The models discussed earlier, in particular Jovanovic and MacDonald (1994a), suggest that the magnitude of change in the moments should depend on the direction and magnitude of the change in the number of firms during a life-cycle phase, which, in turn, depends on industry-specific fundamentals such as the rate of innovation and the difference between the scales of high-tech and low-tech firms. Since most of these fundamentals are not observable in our data, it is not possible to relate the changes in the moments to them directly. Nevertheless, the magnitude of the change in the number of firms and industry output should reflect the strength of these fundamentals. For instance, we expect to observe a more pronounced change in the moments of the firm size distribution in an industry that experiences a mass exit of firms compared to an industry that loses only a small fraction of firms. In Jovanovic and MacDonald (1994a), for example, a higher rate of adoption of the better technology and a larger gap between the sizes of high-tech and low-tech firms lead to a more severe and faster decline in number of firms, a larger growth rate in output, *as well as* a higher rate of increase in average firm size.

We measure the extent of the life-cycle phase using the percent change in the number of firms, ΔN , during that phase. We relate each Δy to ΔN for a given life-cycle phase using a simple projection of the form

$$\Delta y_i = \alpha + \beta_N \Delta N_i + \varepsilon_i, \tag{6}$$

where i indexes industries and ε_i is a projection error that represents the effect of unobservables.

5 Examples

To illustrate our general analysis, we begin with two detailed examples that are suggestive of the general results. For these examples, we classify plants based on their primary 4-digit industry. The results are very similar if, instead, we classify plants into multiple industries based on all SIC codes in which they produce. The examples suggest that there may be substantial inter-industry regularities in the way firm heterogeneity evolves, and that the evolutions may vary systematically depending on the phase of the life-cycle the industry is experiencing.

5.1 Electronic computers

The evolution of key variables in the Electronic Computers industry (SIC 3571) is shown in Figure 7.²⁰ The number of firms evolved non-monotonically, and parts of phases I and II are both clearly visible, albeit not observed in their entirety. The number of firms peaked around 1982, and a shakeout followed, reducing the number of firms to about half its peak value.²¹ Throughout the sample period output increased exponentially and price fell sharply.²² The distribution of the logarithm (base 10) of firm employment is highly positively skewed.²³ Between 1963 and 1982 (phase I), employment tended to decrease stochastically, while between 1982 and 1997 (phase II), it appears to have exhibited a slight stochastic increase. In contrast, the distribution of output increased stochastically during both phases of the life-cycle. The number of plants per firm also declined steadily throughout the two phases, rebounding slightly after 1992.

²⁰According to the 1987 SIC system, this industry is composed of two related 5-digit products: “Computers (excluding word processors, peripherals and parts)” and “Parts for computers”.

²¹Chandler (2001) reports that the shakeout in the Personal Computers (PC) industry started around 1984. The start of the shakeout observed in the data (1982) is very much in line with the this reported timing of the shakeout in the PC sector, even though the definition of Electronic Computers industry is wider than PC’s. Many of the start-ups, including Sinclair, Osborne and Corona, and established firms, such as Texas Instruments, exited the PC industry during the shakeout. Some of these exiting firms reappeared later in the life-cycle of the industry essentially by putting their labels on computers produced overseas.

²²The output and price data are from the NBER/CES Manufacturing Productivity database. As mentioned earlier, since the NBER/CES data has annual observations, we use the NBER/CES output data to generate graphs of the evolution of the output. For these examples, the pattern of output time-series is similar in the NBER/CES data and the Census of Manufacturers.

²³The figures are nonparametric kernel density estimates based on the raw data. These estimates have the advantage of revealing the shape of the density without restricting it to a parameterized family of distributions, and for this reason are widely used for describing the shape of firm size distribution (see, e.g., Cabral and Mata (2002) and Lotti and Santarelli (2004)). We also employ them to avoid – as we are required to do – revealing too much about the specific firms in the tails of the distribution, as would be the case if we used a histogram.

The estimation method assigns a value to the density at any given point in the support of the distribution by performing a “local” averaging of the observed frequency in neighboring points. How many points each neighborhood should include, and how much weight should be put on each, are the main ingredients of kernel estimation. For the kernel estimate of the density of the logarithm of firm size, we used the commonly employed Parzen density estimate

$$\hat{f}_h(x) = \frac{1}{n} \sum_{i=1}^n \frac{1}{h} K\left(\frac{x - \tilde{x}_i}{h}\right),$$

where x is the point where the density is evaluated, $\{\tilde{x}_i\}_{i=1}^n$ are the *logarithms* (base 10) of the observed firm sizes, $K(\cdot)$ is the kernel that determines the nature of “smoothing”, and h is the bandwidth that determines the range over which smoothing applies. We chose a Gaussian kernel of the form

$$K(u) = \frac{1}{\sqrt{2\pi}} e^{-\frac{u^2}{2}}.$$

The value of the bandwidth was selected to be somewhat larger than the optimal bandwidth size, i.e., that which minimizes the mean squared error of $\hat{f}_h(x)$. This extra smoothing ensures that no individual firm sizes are revealed, especially at the tails of the estimated density.

Figure 8 provides more detail on the evolution of employment. While the employment distribution mostly shifted to the left during the 1963-1982 period, it tended to shift to the right between 1982 and 1997, i.e., after the number of firms reached its peak. The overall shift was not monotonic, however, as indicated by the rightward shift between 1982 and 1987, followed by a leftward shift thereafter.

The evolution of the moments of firm employment (in levels) shown in Figure 9 reveals an interesting asymmetry between the two phases of the life-cycle. Before 1982, the mean, the median, and the standard deviation of firm employment declined from their 1963 levels, while the coefficient of variation, the skewness, and the kurtosis increased. After 1982, all these trends appear to have reversed. In other words, during phase I of the life-cycle firm employment became increasingly dominated by smaller firms, more dispersed relative to its mean, and had an increasingly fat right tail, whereas during phase II of the life-cycle it became more symmetric, less dispersed relative to its mean, and eventually had a thinner tail compared to the peak year 1982.

The distribution of the logarithm (base 10) of firm output exhibits an almost monotonic rightward shift during both phases, as shown in Figure 10. The evolution of the moments of firm output (in levels) is shown in Figure 11. Unlike in the case of firm employment, the median and the standard deviation of firm output increased steadily beginning in 1963, and while the mean initially declined slightly, it overall exhibited a strong upward trend. Average firm output in 1997 was about 13 times its value in 1963, and the standard deviation of output actually increased by about 18 times. The behavior of the higher moments, though, is remarkably similar to those of firm employment. The coefficient of variation, skewness, and kurtosis all increased during phase I, peaking in 1982 and thereafter declined till the end of the sample period.

The evolutions of the computer industry and the automobile tire industry discussed in Appendix A look alike in many ways, especially in terms of the time-paths of the number of firms, output, and price. A comparison of the evolutions of the estimated moments of firm output in Figure 2 and in Figure 11 also reveals substantial similarity. In both cases, average firm size tended to rise over time. The standard deviation exhibited monotonic increase in both industries along most of the life-cycle, with exception that at the end of the evolution of the automobile industry, the estimated standard deviation from the model vanishes as an artifact of the assumption that there are only two firm types. The coefficient of variation, skewness, and kurtosis all moved in the positive direction initially and then reversed their trends. The resemblance of the patterns exhibited by two such different industries going through similar life-cycle phases suggests that we might find significant regularity among other industries.

5.2 Semiconductors

The semiconductor industry is our second example; its evolution is summarized in Figure 12.²⁴ Like in the computer industry, output increased substantially over time and price declined. Unlike in the computer industry, however, the number of firms exhibited only an upward trend and no shakeout was observed by 1997. The employment distribution initially shifted to the left between 1963 and 1982, and did not change substantially between 1982 and 1997. The distribution of output changed somewhat between 1963 and 1982, and it exhibited some growth in the right tail and a shortening of the left tail. A much more pronounced shift occurred between 1982 and 1997, during which output increased stochastically. The number of plants per firm declined until roughly 1977, and then rebounded slightly before declining again.

The evolution of the moments of employment shown in Figure 13 is remarkably similar to phase I of the computer industry’s life-cycle. Just as in the computer industry, the mean, the median, and the standard deviation declined, and the higher moments increased, although the magnitudes of change in the moments were different across the two industries. Similarly, the moments of firm output distribution evolved in a way qualitatively similar to phase I of the computer industry. In fact, in the last year phase I was observed in both industries, all the moments of firm output were higher than their starting values. Some moments grew substantially, notably the mean, the standard deviation, and the coefficient of variation, while others experienced more moderate growth. Note also that no reversion occurred in the trends exhibited by the moments in the semiconductor industry, presumably, unlike the computer industry, it had not experienced a shakeout phase between 1963 and 1997. These examples, while only suggestive, raise the possibility of systematic differences in the behavior of the moments during different phases of the life-cycle.

We now turn to the full sample of industries and implement the statistical methodology described above.

6 Results

6.1 Evolution of moments

We begin with the primary-SIC-code-based classification of plants and firms to industries. Implementing the methodology described in Section 4.1, we classified all 322 4-digit industries into four groups based on the pattern the number of firms exhibited: 140 industries showed general growth in the number of firms; 119 industries showed general decline; 53 industries showed a relatively stable pattern or no obvious trend; and 10 industries displayed a clear non-monotonic path, i.e. first increasing and then decreasing number of firms. Note that all of these patterns are consistent with what we

²⁴The 1987 SIC system includes four closely related 5-digit product classes under the definition of this industry: “Integrated microcircuits (including semiconductor networks, micro processors, and MOS memories)”, “Transistors”, “Diodes and rectifiers”, and “Other semiconductor devices (including semiconductor parts, such as wafers and heat sinks)”.

expect to observe if life-cycle effects are present in the data. For instance, a V-shaped time path for the number of firms would be at odds with the general life-cycle pattern for the number of firms; we did not observe this pattern in any of the industries.

In the growing industries, between 1963 and 1997, the number of firms grew, on average by 170.3%; the corresponding figure for declining industries is -47.8% . In Industries with industries with stable firm numbers grew slightly (on average, $\sim 12\%$ increase in the number of firms). The industries showing non-monotonic evolutions experienced, on average, about 60% growth in firm numbers. In terms of output, 270 industries showed positive growth; 32 displayed persistent decline in output; and 20 experienced an inverted U-shaped path. The inverted U-shaped evolution of output is different from the typical monotonic behavior of output over the life-cycle. But these cases only account for only 6% of all industries; presumably these industries are those that entered their final decline during the sample period.

We now classify industries into life-cycle phases based on the evolution of both firm numbers and output. As shown in Table 3, we were able to assign a total of 284 industries – nearly 90% of the industries we study – into four major groups consistent with industry life-cycles: 127 industries exhibited both growing number of firms and output (phase I); 91 industries experienced a decline in number of firms but a growth in output (phase II); 44 industries exhibited no obvious trend in number of firms, but experienced growth in output (phase 3); and 22 industries declined in both measures (decline phase).²⁵

Some general features of these groups of industries are as follows. In industries experiencing phase I, total output increased, on average, by a factor of 20, whereas average output grew by about half that much for firms experiencing phase II, and a little more than one tenth as much for the group experiencing phase III. For declining industries, average output at the end of the sample was about half its initial value. Based on the real price series for each industry (construction described in Appendix B.2.1), industries in phase I saw the average real price decrease by about 25%. The corresponding numbers for phases II and III, and decline, are 30%, 11%, and 20%. These observations are consistent with both life-cycle theory and the trends discovered by Gort and Klepper (1982) and their followers.

Now consider the evolution of moments of the firm size distribution; refer to Table 3. The moments of the firm size distribution evolve very differently in phase I industries in comparison to phase II. Take employment first. As seen in Table 3a, on average, employees per firm declined in phase I, but

²⁵The 38 industries that did not fall into the four major groups can be divided into two subsets – industries whose behavior is consistent with industry life-cycles, but for which there are too few industries to perform much analysis, and others. The first group, 17 industries total, includes an inverted-U evolution of firm numbers alongside growing output (8 industries); a declining number of firms and inverted-U output (5 industries, entering their final decline), and an inverted-U path for firm numbers and declining output (4 industries, also entering their final decline). The remaining 21 industries, just over 6% of those on which we have, display behavior not easily accommodated by existing life cycle models: growing firm numbers of firms and inverted-U output (10 industries); stable number of firms and declining output (5 industries); inverted-U number of firms and output (1 industry); growing number of firms and declining output (2 industries); and stable number of firms and inverted-U output (3 industries).

increased in phase II. The dispersion of firm employment, as measured by the coefficient of variation, increased in phase I by about 33%, but did not change significantly on average in phase II. Skewness and kurtosis both grew during phase I, and shrank in phase II. In phase III, the mean, median and standard deviation did not change significantly, although the higher moments exhibited some positive growth, similar to phase I. In declining industries, average firm employment did not change substantially, and the only significant trends are in skewness and kurtosis, both of which decreased on average, indicating increasingly symmetric and thinner-tailed employment distributions as decline progresses.

Turning to output, Table 3b shows that average and median output grew in all phases, but the rates of growth differed by phase – the highest growth rate occurs in phase II, followed by phase I. While average output grew by roughly the same amount in phase III and in the decline phase, median output grew about 6 times more in declining industries. Growth in the standard deviation was greatest in phase I industries, followed by those in phase II, phase III, and decline. The coefficient of variation increased substantially during phase I and phase III, but did not change by much in industries experiencing phase II or decline. The similarity of the changes in skewness and kurtosis in phases I and III, and in phase II and the decline phase, is notable. In phases I and III, both moments tend to increase, although much more in absolute value for phase I. In phase II and the decline phase, both moments decrease, but at a higher rate in the decline phase.

While the differences between phases seem striking, there is a lot of inter-industry variability in the moments, so whether the inter-phase differences are as clear as they seem is debatable. To get at this, Table 3c compares the trends in the moments for phases I and II. The difference in average growth rates of moments (i.e., phase I minus phase II) are significant except for the mean and the standard deviation of output. Although the magnitudes of the differences in these two moments are quite large, both display substantial inter-industry variability, so their averages are not measured with great precision.

Tables 4 repeats the analysis of the evolution of moments when plants are classified into industries based on all the products they produce. We re-classified industries into the same groups as in Table 3. The re-classification resulted in some change in the number of industries falling into different life-cycle phases, especially phases II and III. Differences of this sort are not surprising since the primary output classification is likely to be a grosser characterization of firm activities in some industries compared to others. Nevertheless, comparing Tables 3 and 4 reveals a qualitatively similar pattern across the two classification schemes. The absolute values and the significance of the growth rates differ across the comparable panels, but the directions of shifts observed in the moments and the relative magnitudes of the growth rates in moments for different phases are generally invariant to the classification scheme. The most notable discrepancy occurs for industries in phase III. In the new classification, during phase III the mean and standard deviation of firm employment grew, while in the primary-SIC-code-based classification these moments changed little.

6.2 Implied evolution of firm size

The evolution of the moments of the size distribution documented in the previous section shows that the size distribution does change in a systematic manner. In this section we provide another way to look at this evolution that makes both its character and economic significance more obvious. We consider an arbitrary, but reasonably descriptive, initial distribution, and allow its moments to evolve in a manner similar to the way the moments were estimated to evolve in the previous section. Specifically, assume the initial distribution of firm size, X , is such that $\ln(\ln X)$ is normally distributed. The computer and semiconductor industries seem to have this feature earlier in their life-cycles, and previous work – especially Cabral and Mata (2001) – also suggests that even the natural log of firm size tends to be positively skewed at least for young cohorts of firms. Then, to depict the “typical” evolution of firm size, we applied the observed average growth rates in the moments of firm size X to this initial distribution successively for each phase of the life-cycle. To do this we used the average percentage growth rates in the moments of firm size in Table 3, effectively assuming each phase has duration 35 years; results based on the statistics from Table 4 are very similar. This process produces new moments corresponding to the completion of each phase, but for a size distribution that is not ln-ln normal. To depict this new distribution we fit a mixture of two ln-normal distributions by choosing the weight given to each, and the moments of the two ln-normal distributions, so that the fitted distribution matches the observed average moments as closely as possible.²⁶

The fitted distributions starting from an initial distribution of firm output and employment are given in Figures 15 and 16, respectively. In both cases, the initial distributions are ln-ln normal with mean $\mu = 0$ and variance $\sigma^2 = 0.2$. All distributions are plotted in the ln-domain. The figures make clear that starting from this initial firm size density, the changes experienced in the moments of firm size along the life-cycle translate into important shifts in the density of firm size. Consistent with the magnitude of the statistics in Tables 3 and 4, the density of firm output experiences more dramatic changes in comparison to the density of employment. And as the figure makes obvious, the evolution of the distributions involves not just a shift in their means or medians, but also significant change in the overall shape of the distribution.

²⁶Specifically, let f_1 and f_2 be two ln-normal densities with corresponding parameters (μ_1, σ_1^2) and (μ_2, σ_2^2) . Consider the mixture density $\alpha f_1 + (1 - \alpha)f_2$, where $\alpha \in (0, 1)$. The set of parameters that define the mixture density $\theta = \{\alpha, \mu_1, \sigma_1^2, \mu_2, \sigma_2^2\}$ are chosen such that the mean, the variance, the skewness, the kurtosis, and the median of the mixture distribution are matched as closely as possible with their counterparts in the data for each phase. For each phase, this involves solving a non-linear system of 5 moment-matching equations with 5 unknowns corresponding to the elements of θ , subject to the constraints that $\alpha \in (0, 1)$ and $\sigma_i > 0$, $i = 1, 2$. We employed a constrained least squares optimization routine written in GAUSS to solve for the parameter set θ that provides the best fit. This process was repeated for each phase to plot the entire evolution of firm size density starting from an initial density.

6.3 Stochastic trends

The analysis of the evolution of the firm size distribution documented above yields three general conclusions: the size distributions are evolving; they evolve in a way that depends on the phase of the industry life-cycle; and the distributions of employment and output evolve differently, with the former presumably blending the employment-increasing effects of growth in output with the employment-reducing effects of growing firm efficiency. In this section we explore these ideas further.²⁷

Refer to Table 5, in which employment results appear in panel (a) and output results in panel (b). The column labelled “Change” provides the percentage of industries within a group that exhibited a statistically significant change in the firm size distribution based on the test statistic D in (3). The columns labelled “Increase” and “Decrease” give the percentage of industries which experienced a statistically significant stochastic increase and decrease, respectively, in firm size, based on the convention we adopted using the test statistics D^- and D^+ defined in (4) and (5).

First, consistent with the evolution of the densities we explored in the previous section, in most industries the size distribution does display a statistically significant evolution. When output is the measure of size, roughly nine of every ten industries in phase II (during which demand continues to grow while firm numbers fall) have a size distribution that is significantly different from the 1963 distribution. Even in declining industries, irrespective of how size is measured, roughly two-thirds of industries display a size distribution that is significantly different from the 1963 distribution.

Second, the evolution is much different depending on both life-cycle phase and how size is measured. With output as the measure of size, 89% phase II industries display an increasing size distribution, whereas between 36.3% and 45.4% of industries in the other phases do so. Indeed, a decreasing distribution is almost as common (between 27.2% and 37.0%) as an increasing distribution in the non-phase II industries. Measuring size by employment, however, increasing size distributions are again most common among phase II industries, but just 48.3% of such industries display an increasing distribution; between 18.1% and 20.0% of non-phase II industries display an increasing distribution. In the non-phase II industries, between 59% and 65.9% of industries display a decreasing size distribution.²⁸ Roughly, in these data, an industry selected at random would be most likely to display an increasing size distribution when size is measured by output, but a decreasing distribution with size measured by employment.

Table 6 repeats the analysis in Table 5 using the alternative classification scheme based on all the SIC industries in which a firm produces. There are some minor differences, but the general conclusions

²⁷A fourth general conclusion is that the evolutions are “complex” in the sense that higher moments generally evolve nonmonotonically, whereas lower moments tend to evolve monotonically. The methods we employ in this section do not shed any further light on this result.

²⁸It would be a mistake to interpret these results too strongly. Recall that our procedure for classifying distributions as, e.g., increasing involves a “large” value for D^+ and a “small” value for D^- . While this is consistent with F_T dominating F_1 in the first order sense, this is not guaranteed. And in fact, whenever D^- is positive, the empirical counterparts to F_T and F_1 do not have F_T dominating F_1 .

based on Table 5 also follow from Table 6.²⁹

6.4 The impact of more dramatic life-cycle phases

Some industries experience more dramatic life-cycle evolutions than others. For example, in Gort and Klepper’s (1982) well-known Table 4, the shakeout leaves firm numbers at just 23% of the peak value in the tire industry, whereas the comparable figure for both the outboard motor and television apparatus-parts industries is 81%. Our data contain many similar contrasts. This raises an interesting question: Do industries experiencing more pronounced life-cycle phases also exhibit more dramatic changes in the firm size distribution?

To explore this issue we employ the simple regression framework in (6). Referring back to Table 3, when size is measured by employment, industries experiencing growth in the number of firms display declining mean employment, whereas industries in which firm numbers are shrinking experience increasing mean employment. Thus, for mean employment, we expect β_N to be negative. By the same logic, β_N should be negative for the median and standard deviation of employment, and positive for the coefficient of variation, skewness and kurtosis. The predicted effects are the same when size is measured by output, except that the coefficient should be positive for the standard deviation.

Theory suggests that when the number of firms is growing, i.e., phase I, firms might react differently to a more rapid evolution in comparison to what they would do in response to a more rapid decline, i.e., during phase II. The idea is that phase I is going to be followed by phase II, which puts firms in a different situation in comparison to phase II, which will be followed by the less evolutionary phase III. To allow for this, we divide industries into those with growing number of firms (industries for which $\Delta N > 0$) and those with declining number of firms (industries for which $\Delta N < 0$), regardless of the pattern of the output. Also, since the dependent variable is a growth rate, there is ample opportunity for “outliers” to find their way into the data, especially for growing industries, where the base for the growth calculation may be small. We report both OLS and least absolute deviation (LAD) regressions.

Table 7 contains the estimates for the growing firm numbers ($\Delta N > 0$) and declining firm numbers ($\Delta N < 0$) industries, focusing on the primary SIC code classification. Interestingly, the estimated effects for the growth in mean, median, the coefficient of variation, skewness and kurtosis are almost uniformly consistent with the anticipated results. The exceptions are the OLS results for mean and median output growth, where the estimated β_N is positive. Given the negative LAD coefficients, it appears that the OLS results are driven by outliers. The estimated coefficients when growth in standard deviation is the dependent variable are of mixed signs, but all are insignificant.

²⁹The patterns of stochastic dominance can help the identification of life cycle phases. For instance, suppose that a stochastic increase in firm output is observed in some industry. Assuming the industry is either in decline or in one of phases I, II, or III, and applying Bayes’ rule on the statistics reported, e.g., in Table 6b, the probability that the industry is in Phase II is about 0.40, whereas the probability that the industry is in decline is about 0.07. If, instead, a stochastic decline in output is observed, the probability that the industry is in Phase II is only 0.08, and the probability that the industry is in decline is about 0.23.

As before, we repeated the analysis in Table 7 using the alternative classification of plants based on both primary and other SIC codes of production. The results, shown in Table 8, are qualitatively similar. In the one case in which the standard deviation is the dependent variable, and the coefficient is significant, it has the expected sign.

Overall, the magnitude of the change in the moments of firm size distribution appears to be significantly associated with the severity of the life-cycle effects.

6.5 Robustness and extensions

In this section, we consider two robustness checks and two extensions.

6.5.1 Small firms

Many manufacturing industries contain a large proportion of small firms, many of which are single-employee firms. As noted by Dunne, Roberts, and Samuelson (1988), potential errors in reporting and classification in the Census of Manufactures are likely to be most common for these firms. Thus, while they account for a minuscule proportion of industry employment and output, these firms may influence the analysis of the moments of firm size, stochastic trends, etc. We repeated the entire analysis with single-employee firms excluded. Our results were not substantially different, and the main conclusions remain.

6.5.2 Aggregate trends

Another important issue is the effect of aggregate economic trends. Two such trends are especially relevant: the general decline in plant employment, and the overall productivity growth in the U.S. economy.

Both trends appear in the Census of Manufactures. Between 1963 and 1997, average firm employment in 4-digit manufacturing industries fell, and average firm output in constant dollars increased. Consider the bias that may be introduced by the first trend. Earlier, we observed that the average firm employment declined on average in industries experiencing a growth in the number of firms, whereas it increased on average in industries that lost firms. If entering firms are increasingly smaller over time as a result of the economy-wide trend in plant employment, then the pattern documented for the industries with escalating number of firms could, in principle, be an artifact of this trend.

To explore the influence of the general trend in plant employment on our results, we normalized each firm's employment by the average employment of a manufacturing firm in a 4-digit industry in a given census year. Note that this normalization changes the mean and the standard deviation of firm employment, but does not affect median, the coefficient of variation, skewness, and kurtosis, by the definitions of these moments. We repeated the entire analysis with this normalization. As might be anticipated, since firms experiencing shrinking employment shrink *relatively* less, and growing firms grow *relatively* more, the main differences were (*i*) a less pronounced decline in the mean and the

median firm employment during phase I, but a more pronounced increase in these moments during phase II, and *(ii)* an increase in the tendency for firm employment to exhibit stochastic increase over time. Overall, however, the main conclusions and the differences observed between phases I and II were not altered substantially.

We followed a similar approach to account for the aggregate trend in firm output, i.e., normalizing firm output (in constant dollars, adjusted for industry price and inflation) by the average of that output across all firms in all 4-digit industries. The results were also qualitatively robust to this normalization. Loosely, the aggregate trends are too small in comparison to the life-cycle effects we emphasize to have a great impact on our findings.

6.5.3 Share of output by firm size class

The description of the changes in firm size distribution so far leaves open the question of what, if any, part of the size distribution accounts for most of the changes that occur, especially in industry output. For example, when an industry experiences phase I, is it small, medium, or large firms that contribute most to output growth? More generally, is there a systematic change in the contribution of firms with different sizes to output along the life-cycle? The theory suggests that the fraction of output accounted by firms of different sizes should change over time. For instance, the model Jovanovic and MacDonald (1994a) studied implies a diminishing importance of small firms in industry output after phase I.

To explore this issue, we defined three firm size classes for each industry at a given point in time: “small firms”, i.e. firms smaller than the 33rd percentile of output, “medium firms”, i.e. firms larger than the 33rd percentile but smaller than the 66th percentile, and “large firms”, i.e. firms larger than the 66th percentile. We then calculated the growth rate in the fraction of total industry output accounted by each of these three size classes between 1963 and 1997 by life-cycle phase. Note that the values of the percentiles we use in dividing firms into size classes also change over time depending on how the size distribution evolves, allowing us to maintain a time-invariant definition of what is small and large. This approach is superior to defining firm size classes based on absolute firm size cut-offs.

The results are in Table 9. Based on the primary-SIC-code-based classification in Table 9a, it appears that the contribution of large firms to output is increasing consistently throughout the life-cycle, and the contribution of small and medium firms declines, especially during phases II and III. When the classification based on all SIC codes of production, as shown in Table 9b, the results change for a few cases, but not substantially. The average growth rate in the fraction of output accounted by small firms is now of the opposite sign and statistically significant, compared to the previous classification. The sign of the average growth rate for medium firms in phase II is also reversed, but it is still insignificant. The contribution of small firms appears to decline on average in all phases, although not significantly in phase III and the decline phase. The contribution of large firms increases, as before, throughout the life-cycle.

Overall, the evidence points to a growing importance of large firms over time, especially during

phases III and the decline phase. The contribution of small firms, on the other hand, appears to diminish over time. Note also that the decline phase does not lead to a particularly significant change in the importance of small firms in industry output.

6.5.4 Number of plants per firm

So far we have focused on firms instead of plants. But the fact that the employment distribution evolves quite differently from the output distribution suggests there may be something of interest at the plant level. For example, output generally grows faster than employment. Is this all due to technology growth, or does substitution of other factors also play a role? Are there significant changes in the composition of single versus multi-plant firms in an industry as its life-cycle unfolds?

We know of no general theory of how the number of plants per firm will change along the life-cycle, but some conjectures follow from existing empirical regularities. Dunne, Roberts and Samuelson (1988) find that most entry and exit is by single-plant firms. Because entry dominates exit in phase I of the life-cycle, the number of plants per firm may decrease if the increase in the number of firms as a result of positive net entry is mostly attributable to single-plant firms.³⁰ If, on the other hand, incumbent firms also open new plants during phase I, then the number of plants per firm can increase. Similarly, because exit dominates entry during phase II, the number of plants per firm can increase. However, if multi-plant firms also shut down some or all of their plants, or if there are mergers and acquisitions that may increase the number of plants owned by merging or acquiring firms, then there can be an increase in the number of plants per firm. Studies of industrial decline also have varying predictions and results on plant closures by single-plant versus multi-plant firms (see e.g., Whinston (1998), Deily (1991)).

We calculated the average growth rate in the number of plants per firm by phase. Table 10 contains the results. Based on the primary-SIC-code-based classification, there is a significant increase in the number of plants per firm on average in phase II, and no other economically or statistically significant changes in plants per firm. However, recall that U.S. manufacturing displays a trend towards fewer plants per firm, as documented by Dunne, Roberts, and Samuelson (1988). To control for this effect, we re-calculated the growth rates after normalizing the number of plants per firm in an industry by the average number of plants per firm across 4-digit industries. The average growth rates for the normalized data indicate a small but significant decline in the number of plants per firm during phase I along with the larger increase in phase II. Under the alternative classification based on all SIC codes, a similar picture emerges.

³⁰Denote the number of firms and plants at the start of a phase by N and N^p , respectively. Also, denote the net increase in the number of firms during a phase by ΔN , and the net increase in the number of plants by ΔN^p . If all the increase in number of firms is due to single-plant firms, i.e. $\Delta N^p = \Delta N$, then

$$\frac{N^p + \Delta N^p}{N + \Delta N} \leq \frac{N^p}{N}, \quad (7)$$

because $N^p \geq N$.

6.6 How did theory perform?

Comparing the predictions summarized in Table 1 with the empirical results leads to three general conclusions. Take, for example, the evolution of the distribution of output in Table 3, panel b. This is the distribution which is conceptually most similar to the distributions in the theory, and so most appropriate for evaluating consistency of theory and data. First, while the models allow declining mean output early in the industry life cycle, they all predict eventual growth in mean output. And the growth in output in the data is dramatic indeed. Given that all the models are based on accumulation of capital of some sort, this consistency is not too surprising. Second, the models generally point to a non-monotonic evolution of asymmetry, with asymmetry mirroring firm numbers. More specifically, growth in firm numbers is generally predicted to increase asymmetry by increasing both skewness and kurtosis. Declining firm numbers are predicted to reduce asymmetry by reducing skewness, but the implications for kurtosis in this case are more model-dependent. Consistent with the theory, we find asymmetry consistently evolving in the same direction as firm numbers. (Recall that even in stage III there is some net increase in firm numbers.)

The models generally suggest that imitative processes tend to drive out heterogeneity after Stage I. While it is the case that the standard deviation of output does grow more slowly as the stages of the life cycle unfold, contrary to what the theory predicts, heterogeneity does grow throughout the life cycle. And this is not simply an artifact of the observation that most firms are a lot larger in the later stages of the life cycle, at least in the sense that the coefficient of variation is also observed to grow throughout the life cycle. One natural explanation for the theories' shortcoming in this respect is that the simpler versions, e.g., Jovanovic and MacDonald (1994a) focus on the large innovations that drive firm numbers, and abstract from the many less dramatic process and product innovations that allow firms to continue to grow and prices to continue to fall. Jovanovic and MacDonald (1994b) allow for this sort of ongoing firm growth and discuss how this process might lead to great heterogeneity, even in mature industries. But there is clearly scope for developing models that make clearer predictions about how firms' innovative and imitative activities deliver both the large, early, innovations that drive firm numbers, as well as the less significant, but ongoing, innovations that fuel continued growth in heterogeneity.

For the case of industrial decline, we find a general increase in average firm output, but not much of a change in the coefficient of variation. In an earlier study of chemical products, Lieberman (1990) found that the coefficient of variation of firm capacity declined in products that exhibited the greatest decline. Our findings do not indicate a substantial decline in the coefficient of variation of neither output nor employment, but somewhat increasing symmetry in both firm distributions and thinning tails. Overall, firms did not tend to shrink on average during decline and firm heterogeneity with respect to the average firm did not decrease. The discrepancy in the behavior of the coefficient of variation in Lieberman (1990) versus in our data can be due to the differences in industries analyzed (only chemical industries versus eclectic industries here), industry definition (narrowly defined versus

more aggregated products as in here), and the measure of firm size (capacity versus output and employment as used here).

7 Summary

Industries have life cycles, and even widely differing industries show similar time series patterns in numbers of firms, output, and price. To these well-documented conclusions we add that: the firm size distributions are evolving; they evolve in a way that depends on the phase of the industry life-cycle; the distributions of employment and output evolve differently; and the evolution of the distributions are more dramatic when the life cycle itself evolves more dramatically. We improve on earlier work by analyzing a greater number of industries over a longer time period in a much larger economy. Existing models are generally consistent with our findings, but the data display more heterogeneity in firm size than existing models suggest.

References

- [1] Agarwal, Rajshree (1998) "Evolutionary Trends in Industry Variables", *International Journal of Industrial Organization*, **16**: 511-525.
- [2] Agarwal, Rajshree, and Gort, Michael (1996) "First-mover Advantage and the Speed of Competitive Entry, 1887-1986", *Journal of Law and Economics*, **44**: 161-177.
- [3] Agarwal, Rajshree, and Gort, Michael (2001) "The Evolution of Markets and Entry, Exit, and Survival of Firms", *Review of Economics and Statistics*, **78**: 489-498.
- [4] Axtell, Robert (2001) "Zipf Distribution of U.S. Firm Sizes", *Science*, **292**: 1818-1820.
- [5] Bartelsman, Eric and Gray, Wayne (1996) "The NBER Manufacturing Database", NBER Technical Working Paper 205.
- [6] Cabral, Luis and Mata, Jose (2001) "On the Evolution of the Firm Size Distribution: Facts and Theory", *American Economic Review*, **93**: 1075-1090.
- [7] Chandler, Alfred D. (2001) *Inventing the Electronic Century: The Epic Story of the Consumer Electronics and Computer Industries*, Harvard University Press, Cambridge, Massachusetts.
- [8] Deily, Mary E. (1991) "Exit Strategies and Plant Closing Decisions: The Case of Steel", *Rand Journal of Economics*, **22**: 250-263.
- [9] Dunne, Timothy, Roberts, Mark, and Samuelson, Larry (1988) "Patterns of Firm Entry and Exit in U.S. Manufacturing Industries", *Rand Journal of Economics*, **19**: 495-515.
- [10] Geroski, Paul (1995) "What do We Know about Entry?", *International Journal of Industrial Organization*, **13**: 421-440.
- [11] Ghemawat, Pankaj and Nalebuff, Barry (1985) "Exit", *Rand Journal of Economics*, **16**: 184-194.
- [12] Ghemawat, Pankaj and Nalebuff, Barry (1990) "The Devolution of Declining Industries", *Quarterly Journal of Economics*, **105**: 167-186.
- [13] Gibbons, Jean Dickinson (1971) *Nonparametric Statistical Inference*. McGraw-Hill.
- [14] Gibrat, R. (1931) *Les Inegalites Economiques*, Librairie du Recueil Sirey.
- [15] Gort, Michael and Klepper, Steven (1982) "Time Paths in the Diffusion of Product Innovations", *Economic Journal*, **92**: 630-653.
- [16] Hodrick and Prescott (1997) "Post-war U.S. Business Cycles: An Empirical Investigation", *Journal of Money, Credit and Banking*, **1**: 1-16.

- [17] Hopenhayn, Hugo (1992) “Entry, Exit, and Firm Dynamics in Long Run Equilibrium”, *Econometrica*, **60**: 1127-1150.
- [18] Ijiri, Yuji and Simon, Herbert (1964) “Business Firm Growth and Size”, *American Economic Review*, **54**: 77-89.
- [19] Ijiri, Yuji and Simon, Herbert (1974) “Interpretations of Departures from the Pareto Curve Firm-Size Distributions”, *Journal of Political Economy*, **82**: 315-331.
- [20] Ijiri, Yuji and Simon, Herbert (1977) *Skew Distributions and the Sizes of Business Firms*, North-Holland, New York.
- [21] Jovanovic, Boyan (1982) “Selection and the Evolution of Industry”, *Econometrica*, **50**: 649-670.
- [22] Jovanovic, Boyan and MacDonald, Glenn (1994a) “The Life Cycle of a Competitive Industry”, *Journal of Political Economy*, **102**: 322-347.
- [23] Jovanovic, Boyan and MacDonald, Glenn (1994b) “Competitive Diffusion”, *Journal of Political Economy*, **102**: 24-52.
- [24] Klepper, Steven and Graddy, Elizabeth (1990) “The Evolution of New Industries and the Determinants of Market Structure”, *Rand Journal of Economics*, **21**: 27-44.
- [25] Klepper, Steven (1996) “Entry, Exit, Growth, and Innovation over the Product Life Cycle”, *American Economic Review*, **86**: 562-583.
- [26] Klepper, Steven and Simons, Kenneth (2000) “The Making of an Oligopoly: Firm Survival and Technological Change in the Evolution of the U.S. Tire Industry”, *Journal of Political Economy*, **108**: 728-760.
- [27] Kumar, Krishna, Rajan, Raghuram, and Zingales, Luigi (2002) “What Determines Firm Size?”, Working paper, Graduate Business School, University of Chicago.
- [28] Lieberman, Marvin (1990) “Exit from Declining Industries: “Shakeout” or “Stakeout”?”, *Rand Journal of Economics*, **21**: 538-554.
- [29] Linton, Oliver, Maasoumi, Esfandiar, and Whang, Yoon-Jae (2003) “Consistent Testing for Stochastic Dominance under General Sampling Schemes”, forthcoming in *Review of Economic Studies*.
- [30] Lotti, Francesca, and Santarelli, Enrico (2004) “Industry Dynamics and the Distribution of Firm Sizes: A Nonparametric Approach”, *Southern Economic Journal*, **70**: 443-466.
- [31] Lucas, Robert (1978) “On the Size Distribution of Business Firms”, *Bell Journal of Economics*, **9**: 508-523.

- [32] Ravn, Morten and Uhlig, Harald (2002) “On Adjusting the Hodrick-Prescott Filters for the Frequency of Observations”, *Review of Economics and Statistics*, **84**: 371-375.
- [33] Rossi-Hansberg, Esteban and Wright, Mark L. J. (2004) “Firm Size Dynamics in the Aggregate Economy”, Unpublished manuscript, Stanford University.
- [34] Siegel, Sidney, and Castellan, Jr., N. John (1988) *Nonparametric Statistics for the Behavioral Sciences*. 2nd Edition, McGraw-Hill
- [35] Simon, Herbert and Bonini, Charles (1958) “The Size Distribution of Business Firms”, *American Economic Review*, **48**: 607-617.
- [36] Simons, Kenneth (2001) “Product Market Characteristics and the Industry Life-Cycle”, Manuscript, University of London.
- [37] Sutton, John (1997) “Gibrat’s Legacy”, *Journal of Economic Literature*, **35**:40-59.
- [38] Whinston, M. D. (1988) “Exit with Multiplant Firms”, *Rand Journal of Economics*, **19**: 74-94.

A Derivation of the implications of Jovanovic and MacDonald (1994a)

In this appendix we analyze the evolution of firm size distribution in the industry life-cycle model of Jovanovic and MacDonald (1994a). We provide only a sketch of the parts of the model we need for the discussion.

Consider a competitive industry for a homogenous good. Time is discrete and the horizon is infinite. A number of entrepreneurs with primitive know-how can choose to enter the industry as soon as a technological advance enables them to produce. At time $t = 0$, a “low-tech” production method becomes available for adoption. A fraction β of the potential entrants adopt this technology and become low-tech producers at period $t = 1$. Nothing else happens till some period $\tilde{T} > 1$, at which point a refinement to the existing technology arrives. This “high-tech” refinement reduces total and marginal cost of production, and thus allows for higher output. The profit-maximizing outputs of high-tech and low-tech firms, q_t^H and q_t^L , satisfy $q_t^H > q_t^L$, i.e. better technology implies higher output. The refinement might be adopted by existing low-tech firms. New entrepreneurs with primitive know-how can also enter the industry in hopes of becoming low-tech producers. A fraction r^o of these entrants succeed in becoming low-tech at period $\tilde{T} + 1$, and the rest exit. From $\tilde{T} + 1$ onwards, a constant fraction r of existing low-tech firms become high-tech each period. As more firms become high-tech, the industry price falls and output increases. Falling price can induce exit, depending on the parameters. If exit occurs, it can be of two types: gradual, in which low-tech firms leave the industry starting at some $T' > \tilde{T} + 1$, or catastrophic, in which all low-tech firms exit en masse at T' .

Let N_t be the total number of firms *active* at time t , N_t^L be the number of *active* low-tech firms, N_t^H be the number of *active* high-tech firms, and N_t^o be the number of entrepreneurs with primitive

know-how who, without either low-tech or high-tech know-how, cannot participate in the market. The evolution of the number of firms for each firm type is as follows:

1. **Period** $t = 0$: $N_t = 0$, $N_t^o = N_0^o$, $N_t^L = 0$, $N_t^H = 0$.
2. **Epoch** $1 \leq t < \tilde{T}$: $N_t = \beta N_0^o$, $N_t^o = 0$, $N_t^L = \beta N_0^o$, $N_t^H = 0$.
3. **Period** $t = \tilde{T}$: $N_t^o = N_{\tilde{T}}^o$, $N_t = \beta N_0^o + N_{\tilde{T}}^o$, $N_t^L = \beta N_0^o$, $N_t^H = 0$.
4. **Epoch** $t \geq \tilde{T} + 1$:

Case 1. No exit:

$$\text{For } t = \tilde{T} + 1 : N_t = N_{\tilde{T}}^L + r^o N_{\tilde{T}}^o, N_t^o = 0, N_t^L = (1-r)N_{\tilde{T}}^L + r^o N_{\tilde{T}}^o, N_t^H = r N_{\tilde{T}}^L,$$

$$\text{For } t > \tilde{T} + 1 : N_t = N_{\tilde{T}+1}^o, N_t^o = 0, N_t^L = (1-r)N_{t-1}^L, N_t^H = N_{t-1}^H + r N_{t-1}^L.$$

For $t > \tilde{T} + 1$, the fraction of high-tech firms in the industry is

$$f_t(r, \beta, N_0^o, r^o, N_{\tilde{T}}^o) = 1 - \frac{(1-r)^{t-\tilde{T}} \left[(1-r)\beta N_0^o + r^o N_{\tilde{T}}^o \right]}{\beta N_0^o + r^o N_{\tilde{T}}^o}.$$

Note that f_t is strictly increasing in t , β , N_0^o , $N_{\tilde{T}}^o$, r and r^o .

Case 2. Gradual exit: For some $T' \geq \tilde{T} + 1$, all low-tech firms exit gradually starting at T' . Let e_t denote the mass of exiting firms in period t . Then,

$$\text{For } t < T' : N_t = N_{\tilde{T}}^L + r^o N_{\tilde{T}}^o, N_t^o = 0, N_t^L = (1-r)N_{\tilde{T}}^L + r^o N_{\tilde{T}}^o, N_t^H = r N_{\tilde{T}}^L,$$

$$\text{For } t \geq T' : N_t = N_{T'-1} - \sum_{i=1}^{t-1} e_i, N_t^o = 0, N_t^L = (1-r)N_{t-1}^L - e_t, N_t^H = N_{t-1}^H + r N_{t-1}^L.$$

For $t > T'$, the fraction of high-tech firms in the industry is

$$f'_t(r, \beta, N_0^o, r^o, N_{\tilde{T}}^o) = 1 - \frac{(1-r)^{t-\tilde{T}} \left[(1-r)\beta N_0^o + r^o N_{\tilde{T}}^o \right] - \sum_{i=\tilde{T}+1}^t r^{i-\tilde{T}} e_i}{\beta N_0^o + r^o N_{\tilde{T}}^o - \sum_{i=\tilde{T}+1}^t e_i}.$$

Note that f'_t is strictly increasing in t , β , N_0^o , $N_{\tilde{T}}^o$, r and r^o .

Case 3. Mass exit: For some $T' \geq \tilde{T} + 1$, all low-tech firms exit at T' .

$$\text{For } t < T' : N_t = N_{\tilde{T}}^L + r^o N_{\tilde{T}}^o, N_t^o = 0, N_t^L = (1-r)N_{\tilde{T}}^L + r^o N_{\tilde{T}}^o, N_t^H = r N_{\tilde{T}}^L,$$

$$\text{For } t \geq T' : N_t = N_{t-1}^H, N_t^o = 0, N_t^L = 0, N_t^H = N_{t-1}^H.$$

The fraction of high-tech firms in the industry is 0 for $t \leq T'$ and 1 for $t > T'$.

Table 2 summarizes the evolution of the moments of the size distribution implied by the evolution of the number of firms.

Consider now the direction of change in each moment over time. For brevity, suppose that Case 2 applies; the other cases can be worked out in a similar manner. Even though time is discrete, for notational convenience we use “derivatives” to calculate the rate of change in each moment. First, note that $\frac{df'_t}{dt} > 0$, because the fraction of firms that are high-tech increases over time. Second, $\frac{dq_t^H}{dt} < 0$ and $\frac{dq_t^L}{dt} < 0$ for $t \in \{\tilde{T} + 1, T' - 1\}$, because as total output increases over time, price declines, depressing output of both types of firms until exit starts at time T' . Once exit starts, the rate of exit by low-tech firms is just enough to maintain a constant price, so q_t^L and q_t^H are both constant after T' , i.e. $\frac{dq_t^H}{dt} = 0$ and $\frac{dq_t^L}{dt} = 0$ for $t \geq T'$.³¹

Let $T^* = \min\{t : f(t) = 1/2\}$ be the time at which exactly half of the firms in the industry are high-tech. The time derivatives of the moments are then as follows

$$\frac{d\mu_t}{dt} = \frac{df'_t}{dt}(q_t^H - q_t^L) + f'_t\left(\frac{dq_t^H}{dt} - \frac{dq_t^L}{dt}\right) + \frac{dq_t^L}{dt} \begin{cases} = 0, & \text{for } t \leq \tilde{T} \\ \geq 0, & \text{for } t \in \{\tilde{T} + 1, T' - 1\} \\ > 0, & \text{for } t \geq T' \end{cases} ,$$

$$\frac{d\sigma_t}{dt} = \frac{1}{2}(f'_t(1 - f'_t))^{-1/2}(1 - 2f'_t)(q_t^H - q_t^L)\frac{df'_t}{dt} + (f'_t(1 - f'_t))^{1/2}\left(\frac{dq_t^H}{dt} - \frac{dq_t^L}{dt}\right) \begin{cases} = 0, & \text{for } t \leq \tilde{T} \\ > 0, & \text{for } t \in \{\tilde{T}, T^* - 1\} \\ < 0, & \text{for } t \geq T' \end{cases} ,$$

$$\frac{dcv_t}{dt} = \frac{\frac{d\sigma_t}{dt}\mu_t - \frac{d\mu_t}{dt}\sigma_t}{\mu_t^2} \begin{cases} = 0, & \text{for } t \leq \tilde{T} \\ \geq 0, & \text{for } t \in \{\tilde{T}, T^* - 1\} \\ < 0, & \text{for } t \geq T^* \end{cases} ,$$

$$\frac{d\gamma_t}{dt} = \frac{1}{(f'_t(1 - f'_t))^{1/2}} \left[-1 - \frac{1}{2f'_t(1 - f'_t)} \right] \frac{df'_t}{dt} < 0, \text{ for all } t$$

$$\frac{d\kappa_t}{dt} = \frac{1}{f'_t(1 - f'_t)} \left(\frac{2f'_t - 1}{f'_t(1 - f'_t)} \right) \frac{df'_t}{dt} \begin{cases} = 0, & \text{for } t \leq \tilde{T} \\ < 0, & \text{for } t \in \{\tilde{T}, T^* - 1\} \\ > 0, & \text{for } t \geq T^* \end{cases} .$$

³¹To see this more clearly, suppose that the period profits are given by the general form $\pi_t^i = p_t q_t - c_i(q_t)$, $i = L, H$, where both cost functions are strictly convex, $c_H(q) < c_L(q)$ and $c'_H(q) < c'_L(q)$ for all q . Then, the optimal choice of outputs in a period are $q_t^i = c_i'^{-1}(p_t)$ for $i = L, H$, where $c_i'^{-1}(p)$ is the inverse function of c'_i . The rates of change in outputs over time during the period $\{\tilde{T} + 1, T' - 1\}$ are then related as

$$\frac{dq_t^L}{dt} = \frac{dc_L'^{-1}(p_t)}{dp} \frac{dp_t}{dt} < \frac{dq_t^H}{dt} = \frac{dc_H'^{-1}(p_t)}{dp} \frac{dp_t}{dt},$$

where the first inequality follows from the facts that $\frac{dc_L'^{-1}(p_t)}{dp} < \frac{dc_H'^{-1}(p_t)}{dp}$ and $\frac{dp_t}{dt} < 0$ for $t \in \{\tilde{T} + 1, T' - 1\}$. For $t \geq T'$, $\frac{dq_t^i}{dt} = 0$ for $i = L, H$, because $\frac{dp_t}{dt} = 0$.

Using the derivatives above, the evolution of the moments can be described as follows. Firm size is constant until time $\tilde{T} + 1$, and can increase or decrease initially. But it increases steadily after exit starts. The higher moments of the size distribution all increase from zero to some positive value as soon as the refinement arrives and some firms adopt it. After that, the standard deviation can increase initially, but eventually declines towards zero; the coefficient of variation can increase or decrease initially, but declines gradually towards zero; skewness is positive initially and falls, approaching zero eventually; kurtosis declines but eventually rises then falls towards zero as nearly all firms are high-tech.³²

Application: The U.S. Automobile Tire Industry

A dramatic example of industry life-cycles is given in Figure 6. The data describe the life-cycle of the U.S. Automobile Tire Industry, as analyzed by Jovanovic and MacDonald (1994a). Note the initial increase in the number of firms and the subsequent shakeout, accompanied by increasing output and falling price. Between 1913 and 1973, average firm output increased by a factor of 112. While sales also increased over time, this increase was not monotonic. Average sales were also more volatile compared to output per firm, especially during the initial phases of the life-cycle during which price declined quickly and fluctuated more. Once price stabilized, however, sales started to trace output more closely.

The evolution of the estimated moments of firm size distribution in the industry is given in Figure 2, which was generated using the estimates of the model's parameters based on the parameterization in Jovanovic and MacDonald (1994a). (This parameterization includes the effect on the tire industry of the general productivity growth in the economy.) Each moment is normalized by its maximum value, so that its highest value is 1. The industry's evolution follows the mass exit pattern described by Case 3 above. The estimated adoption probabilities are $\hat{\beta} = 0.0165$, $\hat{r} = 0.0192$, and $\hat{r}^\circ = 0.1141$. A high-tech firm is estimated to be approximately 97 times larger than a low-tech firm. The estimated arrival date of high-tech know-how is $\tilde{T} = 1914$, and the shakeout episode is estimated to take place in the form of a mass exit at around $T' = 1931$. Finally, the general growth rate in productivity of the economy during the period of analysis was estimated to be 2.93% per year. The productivity growth rate leads to an increase (in addition to the one made possible by the refinement) in average firm size at a constant rate throughout the industry's evolution.

As shown in Figure 2, average firm output is initially high, as only a few firms produce in the industry and the price is high. However, average size declines abruptly as the refinement arrives and price falls. It then increases monotonically until the shakeout as more and more firms become high-

³²Observe that the systematic dependence of the fraction of high-tech firms (f_t and f'_t) on the fundamental parameters generates inter-industry differences in the evolutionary path of the firm size distribution. For instance, changes in the mass of entry (N_0° and N_T°) and the rates of innovation (β , r° , and r), leads to changes in the magnitude of change over time in any moment of the size distribution. Note also that the evolution of average firm size and the standard deviation of firm size also depend on the extent the refinement reduces costs.

tech. Between 1914 and 1931, the negative effect of the decline in price on output is more than offset by the increase in the fraction of high-tech firms and the general productivity increase. In 1931, average size increases abruptly as all low-tech firms exit, and it continues to increase at the rate of the general productivity growth. All other moments surge from zero to a positive value as soon as a few firms adopt the refinement in 1914. From that point on, the standard deviation increases steadily until the shakeout, whereas the coefficient of variation, skewness and kurtosis all decline. After the shakeout all higher moments tend towards zero as only high-tech firms remain, and the size distribution becomes degenerate.

B Data

B.1 Census of Manufactures

We use the census data available for all years 1963, and 1967 through 1997 at five year intervals. The data are at the plant level. Each census year contains information on more than 250,000 individual manufacturing plants. For each plant, the census provides total employment and total value of shipments of the plant by 7-digit product category. This information was used to classify the output of each product manufactured by a plant into a 4-digit industry. The plant level information was then aggregated to the firm level by using the firm identification codes for each plant.

B.1.1 Industries

We grouped firms into industries based on the 1987 SIC system. The SIC system underwent three major changes – in 1972, 1987 and 1997 – during our period of analysis.³³ To obtain industry definitions that are consistent over time, we adhere to the Census Bureau’s re-classification of industries based on the 1987 SIC codes. This re-classified data was made available to us by the California Census Research Data Center. However, using this re-classification is not without problems. The re-classification was made by mapping 7-digit product codes onto 1987 SIC codes. While for most products the re-matching of SIC codes over time appears highly reliable, for some it appears much less so. Furthermore, others have found mistakes in coding of various products; see, e.g., Dunne, Roberts and Samuelson (1988) for a discussion of errors for some 4-digit industries. For these reasons, it is neither possible nor desirable to use all of the 459 4-digit industries in the 1987 SIC system.

To maintain data quality, we dropped the 35 4-digit industries that Dunne, Roberts and Samuelson (1988) found to have mistakes in product code matchings for 1963 and 1967. These industries fall into the 2-digit industry groups 37 (Transportation Equipment) and 38 (Instruments and Related Products). Second, we excluded 65 industries with products “not elsewhere classified (n.e.c)”. These industries are identified by the “n.e.c.” abbreviation in their names. They contain products which could not be classified according to the existing product definitions. Since these industries do not

³³Some minor changes also occurred in other years.

necessarily consist of closely related products, we decided to exclude them from the analysis to keep the industry definitions as uniform as possible.³⁴ Third, 19 industries grouped under the 2-digit category 39 (Miscellaneous Manufacturing Industries) were also deleted based on potential product non-uniformity. Finally, we deleted 18 industries that had either one or more years of missing observations, or exhibited abrupt and substantial jumps in their time series for number of firms and/or output during the years of major SIC code revisions. For instance, in some cases we found that an industry lost a very large number of firms in five years following a change in the SIC code, even though the industry exhibited little trend in the total number of firms before or after. These changes do not look like shakeout episodes and tend to coincide with a revision in SIC codes. We confirmed these abrupt changes by using the industry concordances prepared by the Census Bureau. These concordances identify the fraction of value of shipments in an industry that was re-classified into the new SIC codes after a revision in the codes. In most cases, the abrupt changes in the number of firms and/or output occur for those industries for which a large fraction of the value of shipments was reallocated into other industries. We also found that a few industries had one or more years of missing data; we also dropped these industries. Ultimately, we have 322 industries.

We emphasize that the industries excluded from the analysis are unlikely to be a random sample. Instead the sample is likely biased against new, fast-growing and old, fast-declining industries, since these are more likely to have breaks or abrupt changes in their time series due to industry code revisions, or to higher than usual rates of emergence of new products and disappearance of old products.

B.1.2 Measures of firm size

We use both firm employment and output as measures of firm size. Firm employment is simply the total number of employees in all the plants a firm operates in a given industry. Firm output is obtained from firm value of shipments, which is aggregated from plant level shipments. The value of shipments of a firm at census year t is given by

$$s_t = P_t q_t,$$

where P_t is the nominal price in the 4-digit industry, assumed to be common across firms, and q_t is the firm output. The price P_t can be written as

$$P_t = \pi_t P_t^*,$$

where P_t^* is the real price and π_t is the inflation term. Suppose that a price deflator D_t is available for the nominal price, i.e.

$$D_t = \frac{P_t}{P_\tau},$$

³⁴This is also the approach followed by Dunne, Roberts, and Samuelson (1988).

where τ is a fixed census year (1987). Using this deflator, we obtain the value of output for each year t in year τ dollars as follows

$$s_t^* = \frac{s_t}{D_t} = P_\tau q_t.$$

Thus, we can identify firm output q_t using s_t^* up to a constant, P_τ . The evolution of s_t^* over time is identical to that of q_t . The industry price deflator we use comes from the NBER/CES Manufacturing database, described below.

B.2 NBER/CES Manufacturing Productivity Database

The NBER's Manufacturing Productivity Database contains data on several industry level variables for the years 1958-1996 for all 4-digit manufacturing industries. Detailed information on this dataset is available at www.nber.org/nberces/nbprod96.htm (see also Bartelsman and Gray (1996)) A rich set of industry aggregates are available. The only variables we use from this dataset are the industry price deflator (using 1987 as the base year) and the value of shipments. The price deflator is used to obtain industry output in constant dollars as described in the previous section. To avoid any potential problems of conformity, the value of shipments from the NBER/CES data was used only in producing Figure 5 that describes the time-series behavior of output for certain industries, but not in the empirical analysis, which utilizes the census data. The NBER/CES method of calculating the value of shipments for an industry is not a plant-based approach, unlike what we did in our calculations from the LRD. Rather, the NBER/CES data was based on published data from the Census, and changes in industry codes over time were handled by using the published industry concordances in 1972 and 1987. While practical, this approach introduces some noise. Our approach instead relies on the re-classification of individual 7-digit products into industry codes used by the Census Bureau, considering the changes that have taken place in the industry codes in 1972 and 1987. Product value of shipments were then aggregated to the plant level shipments and summed over all plants to yield the industry level shipments. This approach provides more accurate output measures than the NBER/CES dataset.

B.2.1 Industry price series

Using the price deflator, D_t , from the NBER/CES database and the Consumer Price Index (CPI) series, π_t , available from the 2001 Statistical Abstract of the United States, we calculated a real price series for each industry, where the real price of the industry output in 1963 is normalized to 1. In other words, real price in year t is given by

$$P_t^* = \frac{P_t}{\pi_t} = \frac{D_t P_\tau}{\pi_t},$$

for $t \geq 1963$ and $P_{1963}^* = 1$. This calculation assumes that the unmeasured quality increase in a manufactured product was the same as the unmeasured quality change in the CPI commodity bundle.

Model	Distribution	Mean	Std.	C.V.	Skewness	Kurtosis
Jovanovic and MacDonald (1994a)	↑ or ↓, then ↑	↑ or ↓, then ↑	↑ then ↓	↑ then ↓	↑ then ↓	↑ then ↓ then ↑
Jovanovic and MacDonald (1994b)	↑ or ↓	↑	↑ then ↓ ¹	↑ or ↓, then ↓ ¹	↑ then ↓ ¹	↑ then ↓ ¹
Klepper (1996)	↑ or ↓, then ↑	↑ or ↓, then ↑	↑ or ↓, then ↓	↑ or ↓, then ↓	↑ or ↓, then ↓	↑ or ↓, then ↓
Jovanovic (1982)	↑ or ↓, then ↑	↑ or ↓, then ↑	↑ ²	?	?	?

¹ The prediction applies only to a single cycle of innovation and imitation.

² Increase is non-monotonic.

? No or ambiguous prediction

Table 1: Predictions of different models on the evolution of the moments of firm size distribution

Moments of firm size distribution

<i>Time (t)</i>	μ_t	σ_t	cv_t	γ_t	κ_t
$\{1, T\}$	q_1^L	0	0	0	0
$\{T + 1, T' - 1\}$	$(1 - f_t)q_t^L + f_tq_t^H$	$\sqrt{f_t(1 - f_t)}(q_t^H - q_t^L)$	$\frac{\sqrt{f_t(1 - f_t)}(q_t^H - q_t^L)}{(1 - f_t)q_t^H + f_tq_t^L}$	$\frac{1 - 2f_t}{\sqrt{f_t(1 - f_t)}}$	$\frac{1}{(1 - f_t)f_t}$
$\geq T'$ (Case 1)	Same as epoch $T' > t \geq T + 1$				
$\geq T'$ (Case 2)	$(1 - f'_t)q_{T'}^L + f'_tq_{T'}^H$	$\sqrt{f'_t(1 - f'_t)}(q_{T'}^H - q_{T'}^L)$	$\frac{\sqrt{f'_t(1 - f'_t)}(q_{T'}^H - q_{T'}^L)}{(1 - f'_t)q_{T'}^L + f'_tq_{T'}^H}$	$\frac{1 - 2f'_t}{\sqrt{f'_t(1 - f'_t)}}$	$\frac{1}{(1 - f'_t)f'_t}$
$\geq T'$ (Case 3)	$q_{T'}^H$	0	0	0	0

Table 2: The evolution of the moments of firm size distribution in Jovanovic and MacDonald (1994a)

a.

MEASURE OF FIRM SIZE: EMPLOYMENT

Phase	Pattern in		No. of Industries	Average percent growth rate between 1963 and 1997 in						
	No. of firms	Output		No. of firms	Mean	Median	S.D.	C.V.	Skewness	Kurtosis
I	Growing	Growing	127	175.2 [9.2]	-27.6 [-8.1]	-24.6 [-5.8]	-8.31 [-1.7]	33.0 [6.4]	14.9 [7.4]	23.1 [7.9]
II	Declining	Growing	91	-43.2 [19.9]	72.5 [4.9]	66.3 [3.5]	115.7 [2.5]	8.5 [1.3]	-6.6 [3.9]	-9.2 [-4.0]
III	Stable or no trend	Growing	44	17.4 [2.8]	-5.7 [-0.7]	-12.7 [-0.8]	12.0 [1.2]	24.3 [3.4]	9.1 [3.4]	10.4 [3.6]
Decline	Declining	Declining	22	-65.6 [15.2]	-6.3 [-0.5]	3.9 [0.1]	-9.4 [-0.8]	0.4 [0.03]	-13.6 [-3.6]	-18.7 [-3.8]

Note: t-statistics in brackets. Bold indicates significance at 10% or lower levels.

b.

MEASURE OF FIRM SIZE: OUTPUT

Phase	Pattern in		No. of Industries	Average percent growth rate between 1963 and 1997 in						
	No. of firms	Output		No. of firms	Mean	Median	S.D.	C.V.	Skewness	Kurtosis
I	Growing	Growing	127	175.2 [19.0]	191.7 [1.9]	65.0 [2.4]	687.2 [1.7]	51.5 [8.6]	16.9 [8.9]	25.6 [8.9]
II	Declining	Growing	91	-43.2 [3.1]	300.3 [8.5]	272.3 [4.3]	373.7 [5.6]	13.8 [2.2]	-6.1 [-3.6]	-9.1 [4.1]
III	Stable or no trend	Growing	44	17.4 [6.2]	102.5 [5.8]	24.2 [1.2]	182.3 [5.0]	37.1 [4.8]	8.2 [3.1]	10.1 [3.2]
Decline	Declining	Declining	22	-65.6 [4.3]	102.3 [2.4]	132.7 [1.8]	112.3 [2.9]	9.99 [0.9]	-11.1 [-3.0]	-15.7 [-3.4]

Note: t-statistics in brackets. Bold indicates significance at 10% or lower levels.

c.

Difference in the average growth rates of moments (Phase I minus Phase II)

	Mean	Median	S.D.	C.V.	Skewness	Kurtosis
FIRM EMPLOYMENT	-100.1 [-6.6]	-90.9 [-4.6]	-124.0 [-2.6]	24.5 [2.9]	21.5 [8.0]	32.3 [8.5]
FIRM OUTPUT	-108.6 [-1.0]	-207.3 [-3.0]	313.5 [0.7]	37.6 [4.3]	23.0 [8.9]	34.6 [9.3]

Note: t-statistics for two-sample comparison of means with unequal variances are in brackets.

Bold indicates significance at 10% or lower levels.

Table 3. Evolution of the moments by life-cycle phase (primary SIC code classification)

a.										
MEASURE OF FIRM SIZE: EMPLOYMENT										
Phase	Pattern in		No. of Industries	Average percent growth rate between 1963 and 1997 in						
	No. of firms	Output		No. of firms	Mean	Median	S.D.	C.V.	Skewness	Kurtosis
I	Growing	Growing	128	175.7 [8.81]	-15.6 [-3.9]	-21.7 [-5.0]	12.9 [1.9]	38.1 [7.1]	15.6 [7.4]	24.2 [7.7]
II	Declining	Growing	76	-39.0 [17.4]	103.6 [5.3]	109.7 [3.7]	165.0 [2.8]	12.3 [1.4]	-5.8 [-2.9]	-8.3 [-3.0]
III	Stable or no trend	Growing	37	12.4 [1.6]	24.6 [1.9]	3.4 [0.2]	46.1 [2.9]	26.6 [2.1]	5.7 [1.8]	9.6 [2.0]
Decline	Declining	Declining	28	-61.3 [16.0]	-15.6 [-1.2]	-13.3 [25.2]	-11.4 [-1.0]	12.1 [1.2]	-6.2 [-1.8]	-9.5 [-5.3]

Note: t-statistics in brackets. Bold indicates significance at 10% or lower levels.

b.										
MEASURE OF FIRM SIZE: OUTPUT										
Phase	Pattern in		No. of Industries	Average percent growth rate between 1963 and 1997 in						
	No. of firms	Output		No. of firms	Mean	Median	S.D.	C.V.	Skewness	Kurtosis
I	Growing	Growing	128	175.7 [8.81]	413.0 [1.8]	69.7 [2.8]	1302.7 [1.7]	53.8 [8.4]	16.4 [8.0]	24.9 [8.1]
II	Declining	Growing	76	-39.0 [17.4]	381.3 [8.1]	444.8 [3.6]	492.6 [5.6]	19.6 [2.3]	-5.0 [-2.5]	-7.4 [-2.7]
III	Stable or no trend	Growing	37	12.4 [1.6]	183.4 [6.0]	81.1 [3.4]	324.0 [4.2]	48.3 [3.0]	7.8 [2.5]	12.3 [2.6]
Decline	Declining	Declining	28	-61.3 [16.0]	78.0 [1.9]	69.6 [1.4]	89.2 [2.7]	14.9 [1.6]	-8.2 [-2.6]	-12.2 [2.9]

Note: t-statistics in brackets. Bold indicates significance at 10% or lower levels.

c.						
Difference in the average growth rates of moments (Phase I minus Phase II)						
	Mean	Median	S.D.	C.V.	Skewness	Kurtosis
FIRM EMPLOYMENT	-119.2 [-6.0]	-131.5 [-4.4]	-152.1 [-2.5]	25.7 [2.6]	21.4 [7.3]	32.6 [7.8]
FIRM OUTPUT	31.7 [0.1]	-375.1 [-2.9]	810.2 [1.1]	34.2 [3.2]	21.5 [7.5]	32.4 [7.9]

Notes: t-statistics for two-sample comparison of means with unequal variances are in brackets. Bold indicates significance at 10% or lower levels.

Table 4. Evolution of the moments by life cycle phase (primary and other SIC codes)

a.

MEASURE OF FIRM SIZE: EMPLOYMENT						
Phase	Pattern in		No. of Industries	Percent of industries with stochastic		
	Number of firms	Output		Change ^A	Increase ^B	Decrease ^C
I	Growing	Growing	127	71.6	19.6	62.9
II	Declining	Growing	91	75.8	48.3	38.4
III	Stable/No trend	Growing	44	79.5	20.0	65.9
Decline	Declining	Declining	22	68.2	18.1	59.0

b.

MEASURE OF FIRM SIZE: OUTPUT						
Phase	Pattern in		No. of Industries	Percent of industries with stochastic		
	Number of firms	Output		Change ^A	Increase ^B	Decrease ^C
I	Growing	Growing	127	72.4	44.8	37.0
II	Declining	Growing	91	89.0	86.9	11.4
III	Stable/No trend	Growing	44	61.3	45.4	31.8
Decline	Declining	Declining	22	59.0	36.3	27.2

Notes:

A. Significant at 10% or lower levels based on a **two-sided** Kolmogorov-Smirnov test.

B. Significant at 10% or lower levels based on a **one-sided** Kolmogorov-Smirnov test, following the convention in the text.

C. Significant at 10% or lower levels based on a **one-sided** Kolmogorov-Smirnov test, following the convention in the text.

Table 5. Summary of trends in firm size distribution by life-cycle phase (primary SIC code classification)

a.

MEASURE OF FIRM SIZE: EMPLOYMENT						
Phase	Pattern in		No. of Industries	Percent of industries with stochastic		
	Number of firms	Output		Change ^A	Increase ^B	Decrease ^C
I	Growing	Growing	126	82.5	24.6	74.6
II	Declining	Growing	76	83.3	62.1	54.5
III	Stable/No trend	Growing	37	77.7	33.3	57.5
Decline	Declining	Declining	28	67.8	10.7	67.8

b.

MEASURE OF FIRM SIZE: OUTPUT						
Phase	Pattern in		No. of Industries	Percent of industries with stochastic		
	Number of firms	Output		Change ^A	Increase ^B	Decrease ^C
I	Growing	Growing	126	73.0	50.8	30.9
II	Declining	Growing	76	89.3	87.8	7.5
III	Stable/No trend	Growing	37	70.3	66.6	18.5
Decline	Declining	Declining	28	64.2	35.7	50.0

Notes:

A. Significant at 10% or lower levels based on a **two-sided** Kolmogorov-Smirnov test

B. Significant at 10% or lower levels based on a **one-sided** Kolmogorov-Smirnov test, following the convention in the text.

C. Significant at 10% or lower levels based on a **one-sided** Kolmogorov-Smirnov test, following the convention in the text.

Table 6. Summary of trends in firm size distribution by life-cycle phase (primary and other SIC codes)

Dependent variable:	EMPLOYMENT				OUTPUT			
	Growing number of firms ($\Delta N > 0$)		Declining number of firms ($\Delta N < 0$)		Growing number of firms ($\Delta N > 0$)		Declining number of firms ($\Delta N < 0$)	
Percent growth in	OLS	LAD	OLS	LAD	OLS	LAD	OLS	LAD
Mean	-0.041 [-3.48]	-0.039 [-2.07]	-0.757 [-1.33]	-0.061 [-0.12]	2.359 [1.67]	-0.023 [-2.10]	-3.654 [-2.24]	-1.830 [-1.94]
R ² or Pseudo R ²	0.05	0.02	0.02	0.0005	0.23	0.0004	0.06	0.02
Median	-0.056 [-4.02]	-0.054 [-2.32]	-1.288 [-1.72]	-0.097 [0.26]	0.662 [1.75]	-0.064 [-1.80]	-4.718 [-2.02]	-0.297 [-1.99]
R ² or Pseudo R ²	0.06	0.03	0.03	0.0003	0.24	0.005	0.03	0.01
Standard deviation	0.007 [0.34]	0.002 [0.09]	1.058 [0.66]	0.582 [1.33]	9.712 [1.63]	0.130 [2.17]	-0.982 [-0.38]	-0.318 [-0.25]
R ² or Pseudo R ²	0.0008	0.0002	0.003	0.008	0.22	0.002	0.001	0.0001
Coeff. of variation	0.130 [4.55]	0.144 [10.82]	0.575 [2.56]	0.437 [1.81]	0.165 [4.72]	0.164 [6.59]	0.546 [2.73]	0.458 [1.88]
R ² or Pseudo R ²	0.24	0.15	0.05	0.04	0.29	0.12	0.04	0.02
Skewness	0.060 [5.91]	0.062 [12.47]	0.317 [5.16]	0.284 [3.74]	0.061 [7.37]	0.071 [0.13]	0.297 [4.95]	0.402 [5.87]
R ² or Pseudo R ²	0.35	0.18	0.17	0.09	0.37	0.18	0.16	0.12
Kurtosis	0.087 [6.81]	0.084 [8.20]	0.425 [5.36]	0.404 [4.35]	0.091 [6.98]	0.095 [8.27]	0.413 [5.27]	0.514 [7.05]
R ² or Pseudo R ²	0.33	0.18	0.18	0.09	0.37	0.17	0.18	0.13

Notes: t-statistics in brackets. Bold indicates significance at 10% or lower levels with a two-sided test. For the least absolute deviation (LAD) regressions, Pseudo R² is reported.

Table 7. Growth in the moments of firm size versus industry growth (primary SIC code classification)

Dependent variable:	EMPLOYMENT				OUTPUT			
	Growing number of firms ($\Delta N > 0$)		Declining number of firms ($\Delta N < 0$)		Growing number of firms ($\Delta N > 0$)		Declining number of firms ($\Delta N < 0$)	
	OLS	LAD	OLS	LAD	OLS	LAD	OLS	LAD
Percent growth in								
Mean	-0.038 [-2.04]	-0.044 [-1.97]	-0.761 [-1.06]	-0.213 [0.49]	-0.031 [-1.81]	-0.038 [-1.87]	-3.948 [-2.03]	-1.851 [-2.61]
R ² or Pseudo R ²	0.03	0.02	0.01	0.003	0.22	0.04	0.06	0.01
Median	-0.052 [-3.92]	-0.053 [-2.17]	-1.506 [-1.62]	-0.259 [0.59]	-0.076 [-1.72]	-0.075 [-2.37]	-5.205 [-1.53]	-0.066 [-0.68]
R ² or Pseudo R ²	0.05	0.04	0.03	0.004	0.22	0.006	0.02	0.02
Standard deviation	0.034 [0.88]	0.008 [0.22]	1.289 [0.59]	0.256 [0.68]	0.143 [2.23]	0.148 [2.11]	-1.347 [-0.41]	-0.844 [0.59]
R ² or Pseudo R ²	0.01	0.001	0.005	0.003	0.01	0.002	0.002	0.002
Coeff. of variation	0.146 [6.29]	0.155 [9.18]	0.460 [1.95]	0.397 [1.72]	0.185 [5.36]	0.171 [7.36]	0.522 [2.14]	0.454 [1.93]
R ² or Pseudo R ²	0.29	0.17	0.03	0.03	0.34	0.17	0.03	0.02
Skewness	0.065 [7.47]	0.066 [10.68]	0.222 [3.14]	0.232 [2.48]	0.066 [7.28]	0.081 [8.64]	0.216 [3.17]	0.343 [9.64]
R ² or Pseudo R ²	0.39	0.21	0.08	0.04	0.41	0.18	0.07	0.09
Kurtosis	0.095 [7.42]	0.092 [9.10]	0.305 [3.28]	0.282 [2.44]	0.096 [6.16]	0.091 [6.19]	0.306 [3.53]	0.429 [7.44]
R ² or Pseudo R ²	0.37	0.19	0.09	0.03	0.38	0.16	0.08	0.10

Notes: t-statistics in brackets. Bold indicates significance at 10% or lower levels with a two-sided test. For the least absolute deviation (LAD) regressions, Pseudo R² is reported.

Table 8. Growth in the moments of firm size versus industry growth (primary and other SIC codes)

a. Primary SIC code based classification

Phase	Pattern in		No. of Industries	Average percent growth in the fraction of output accounted by		
	No. of firms	Output		Small firms	Medium Firms	Large Firms
I	Growing	Growing	127	16.6 [0.4]	-7.6 [-0.3]	2.7 [4.7]
II	Declining	Growing	91	-19.6 [-2.7]	-10.5 [-1.4]	3.5 [5.9]
III	Stable or no trend	Growing	44	-10.5 [-3.2]	-6.4 [-2.7]	4.5 [4.0]
Decline	Declining	Declining	22	8.6 [0.3]	-5.2 [-0.3]	4.9 [2.3]

Notes: t-statistics in brackets. Bold indicates significance at 10% or lower levels.

b. Classification based on all SIC codes of production

Phase	Pattern in		No. of Industries	Average percent growth in the fraction of output accounted by		
	No. of firms	Output		Small firms	Medium Firms	Large Firms
I	Growing	Growing	128	-13.8 [-2.6]	-14.6 [-2.3]	3.8 [6.1]
II	Declining	Growing	76	-17.2 [-1.8]	10.5 [0.6]	2.9 [3.6]
III	Stable or no trend	Growing	37	-8.4 [0.4]	-11.5 [2.7]	4.6 [3.9]
Decline	Declining	Declining	28	-6.7 [-0.2]	-12.5 [-1.1]	5.0 [2.9]

Notes: t-statistics in brackets. Bold indicates significance at 10% or lower levels.

Table 9. The change in the contribution of firms with different sizes to industry output by life-cycle phase

a. Primary SIC code based classification

Phase	Pattern in		No. of Industries	Average percent growth in the number of plants per firm	
	No. of firms	Output		Non-normalized	Normalized
I	Growing	Growing	127	0.02 [0.01]	-2.1 [-2.1]
II	Declining	Growing	91	13.0 [7.4]	10.8 [6.3]
III	Stable or no trend	Growing	44	-1.1 [0.4]	-3.0 [-1.1]
Decline	Declining	Declining	22	-1.9 [0.6]	-3.9 [1.3]

Notes: t-statistics in brackets. Bold indicates significance at 10% or lower levels.

b. Classification based on all SIC codes of production

Phase	Pattern in		No. of Industries	Average percent growth in the number of plants per firm	
	No. of firms	Output		Non-normalized	Normalized
I	Growing	Growing	128	0.6 [0.7]	-1.7 [-2.0]
II	Declining	Growing	76	14.5 [6.5]	11.9 [5.3]
III	Stable or no trend	Growing	37	-1.6 [-0.6]	-3.9 [-1.7]
Decline	Declining	Declining	28	-3.6 [-1.5]	-5.6 [-2.4]

Notes: t-statistics in brackets. Bold indicates significance at 10% or lower levels.

Table 10. The change in the number of plants per firm by life-cycle phase

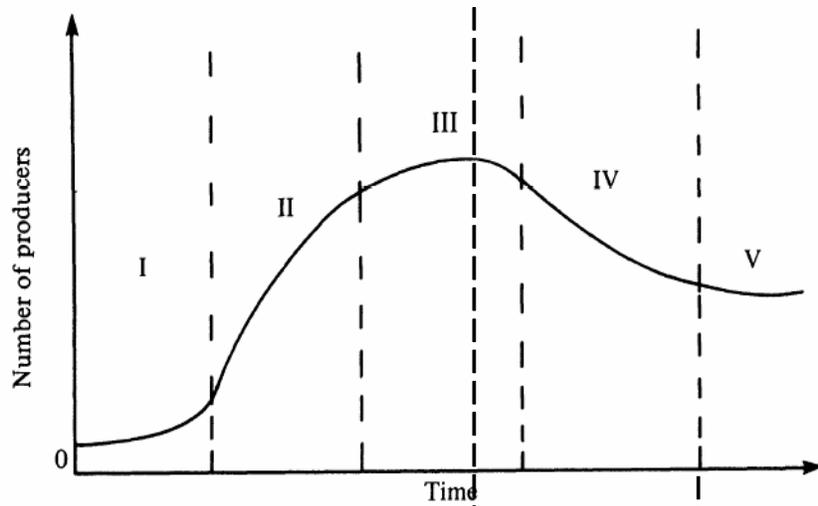


Fig. 1. The five stages of new product industries.

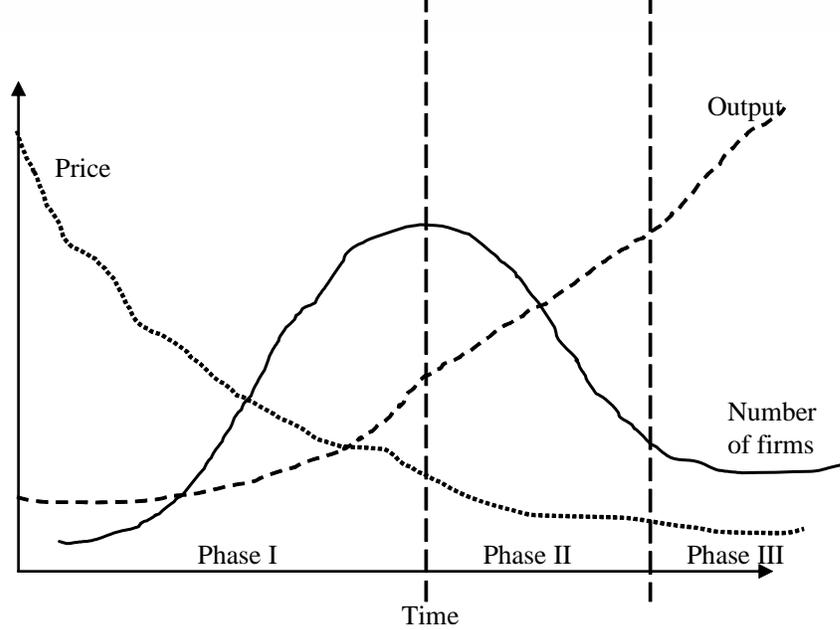


Figure 1: Typical evolution of the number of firms, output, and price in an industry:

Top panel: Gort and Klepper's stages of life-cycle (reproduced from Figure 1 in Gort and Klepper (1982)).

Bottom panel: Three phases of life-cycle used in this study.

Evolution of the estimated moments of firm size in the U.S. Tire Industry

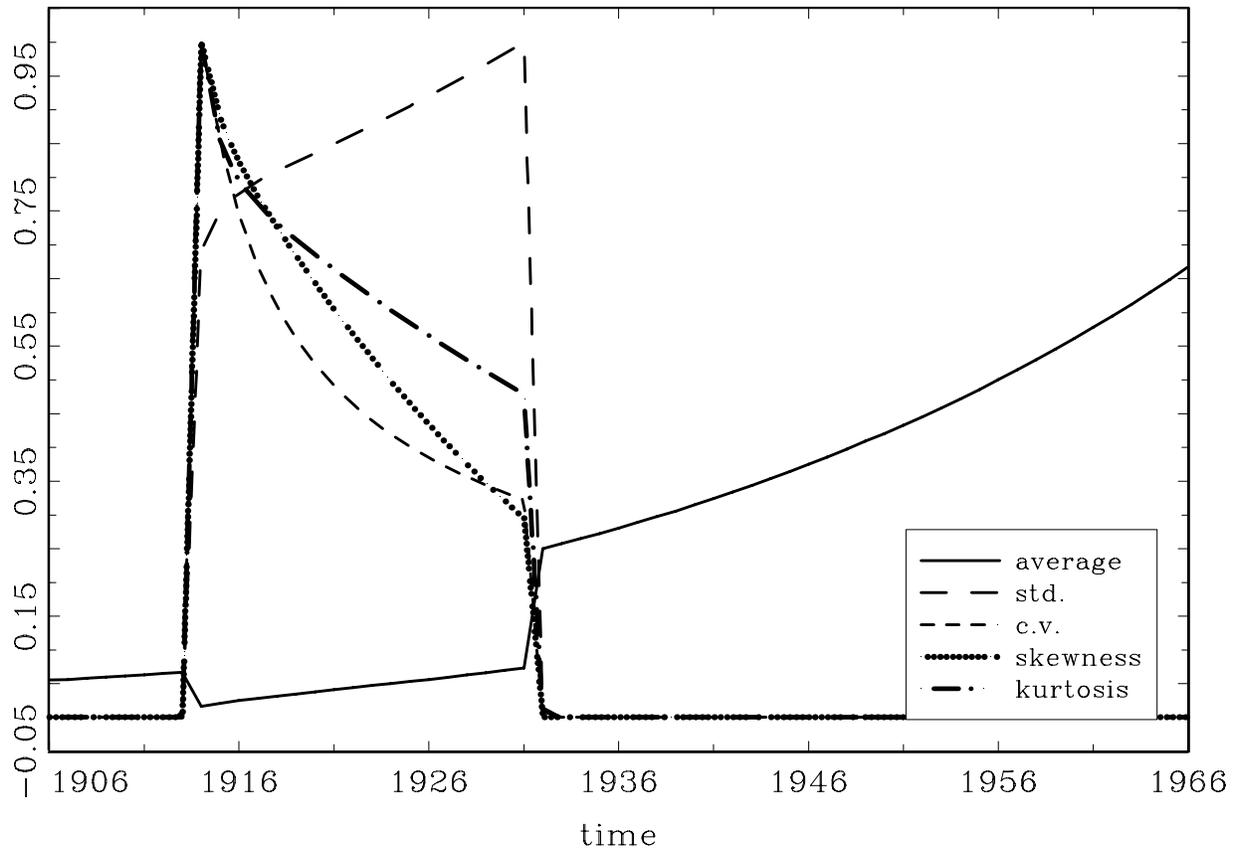


Figure 2: Evolution of the estimated moments of firm output in the U.S. Automobile Tire Industry, 1906-1946. (Based on the estimates from Jovanovic and MacDonald (1994a))

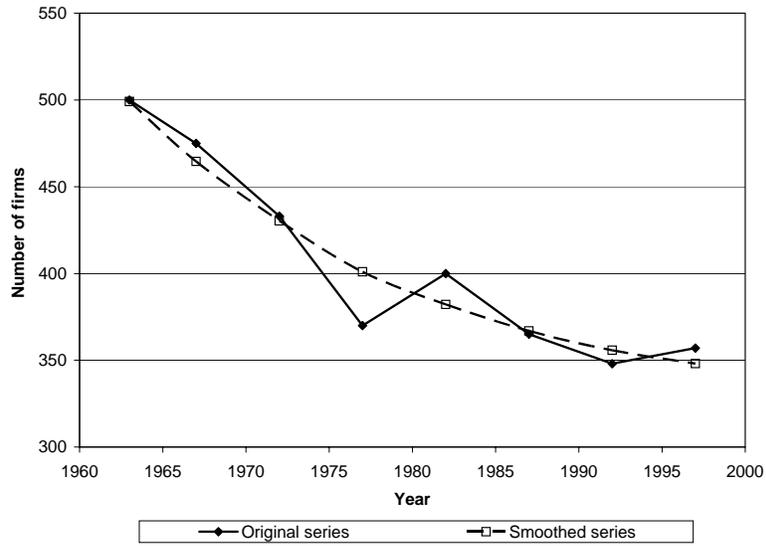


Figure 3: An example of the implementation of the HP-filter for time series data for the number of firms in an industry.

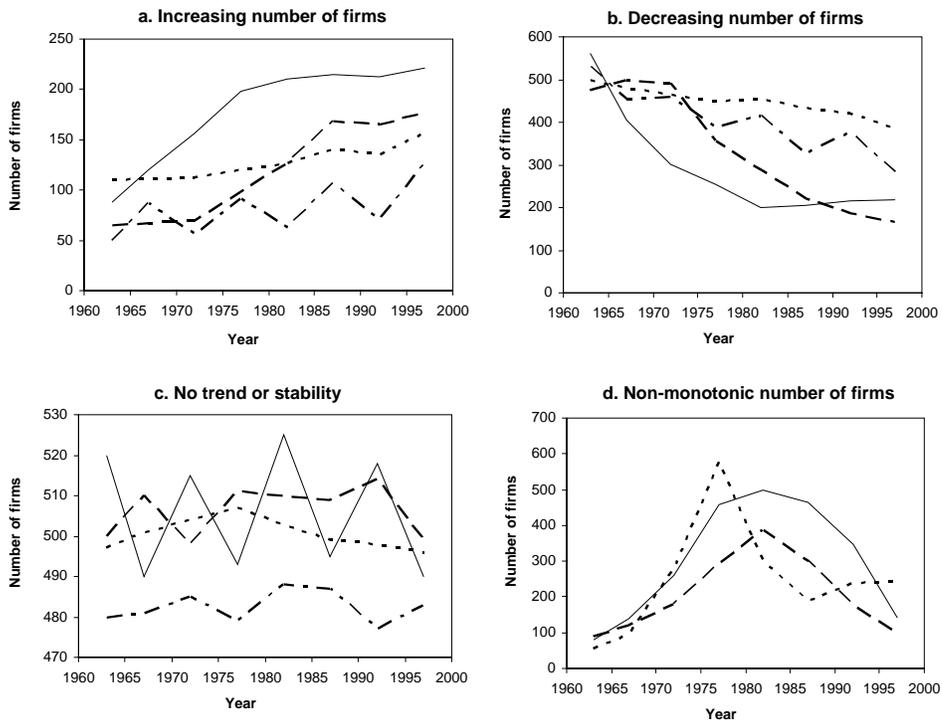


Figure 4: Sample time-paths for number of firms

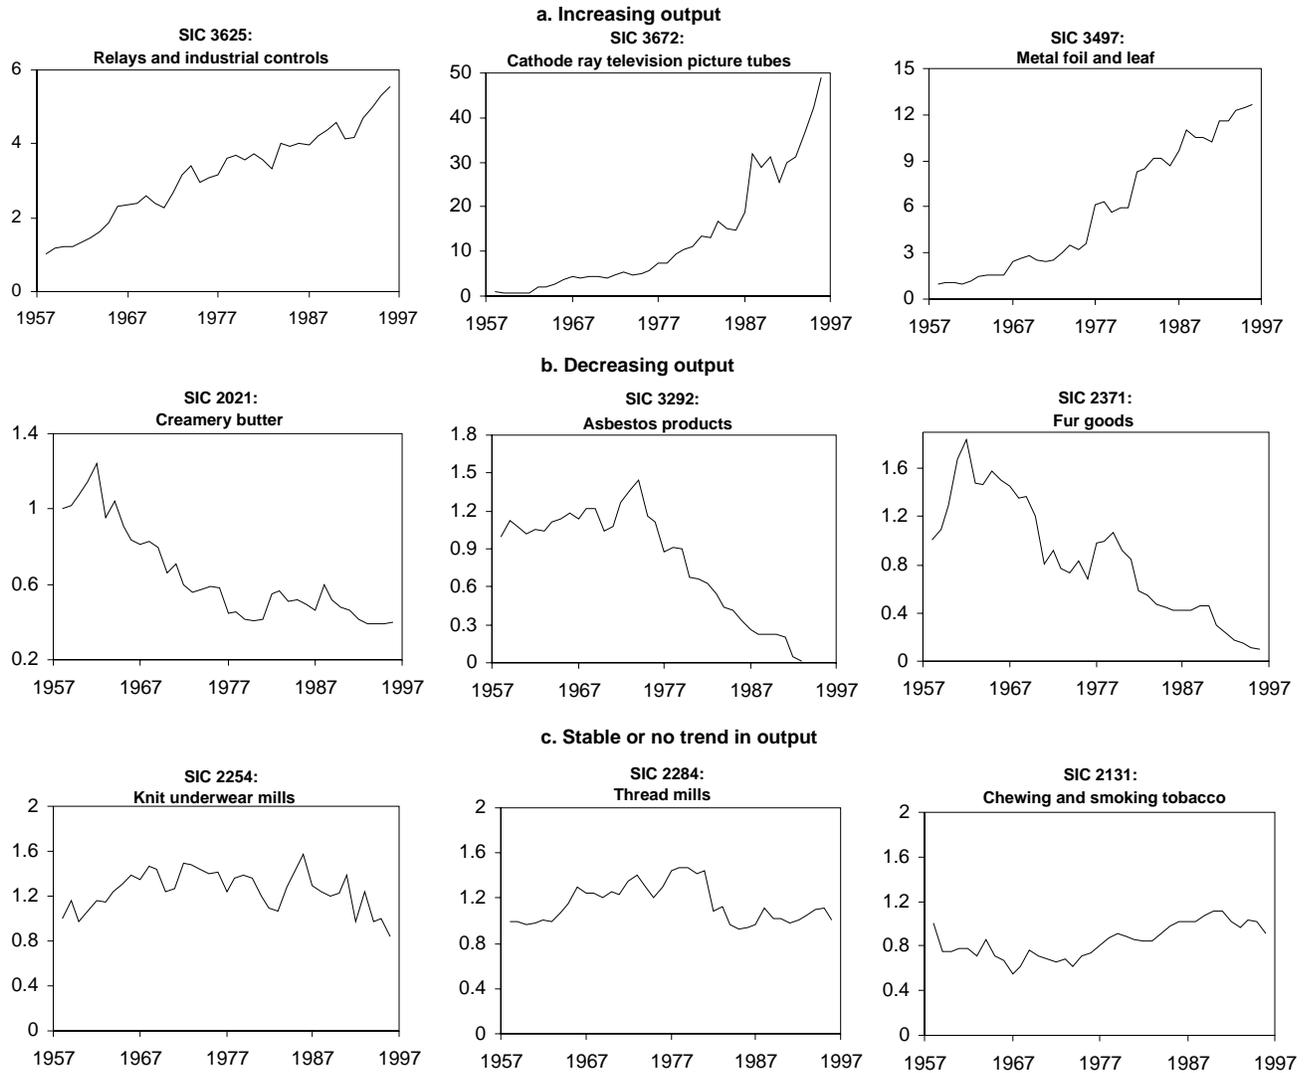


Figure 5: Sample time-paths for output (Output is defined as the value of output in 1987 dollars. Time series is normalized so that the 1958 value is equal to 1.)

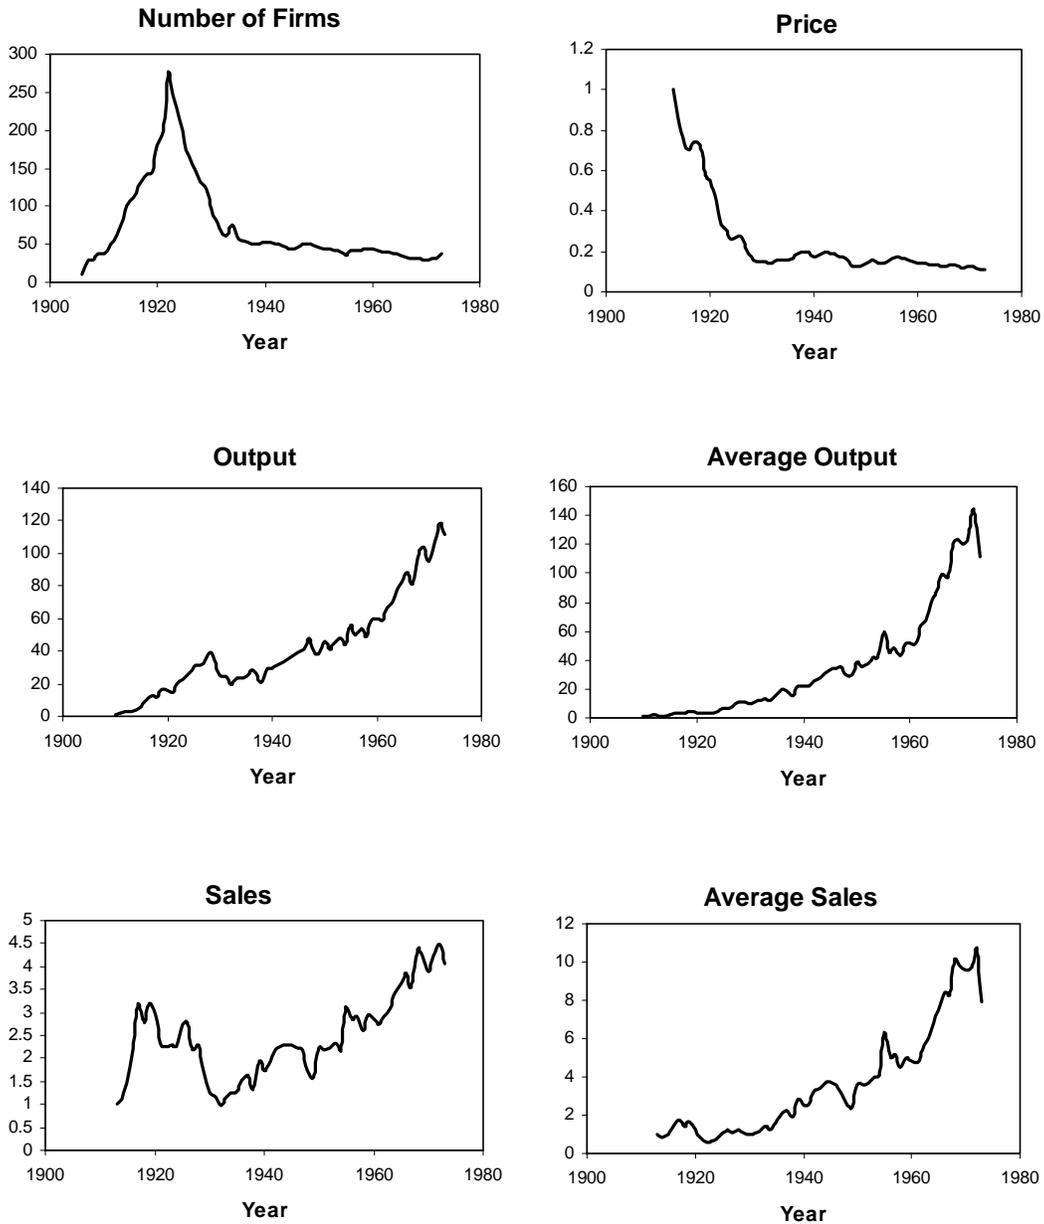


Figure 6: Evolution of key industry aggregates in the U.S. Automobile Tire Industry 1906-1973. (All time series, except for the number of firms, are normalized so that the initial value is 1)

Evolution of Key Variables: SIC 3571 Electronic Computers

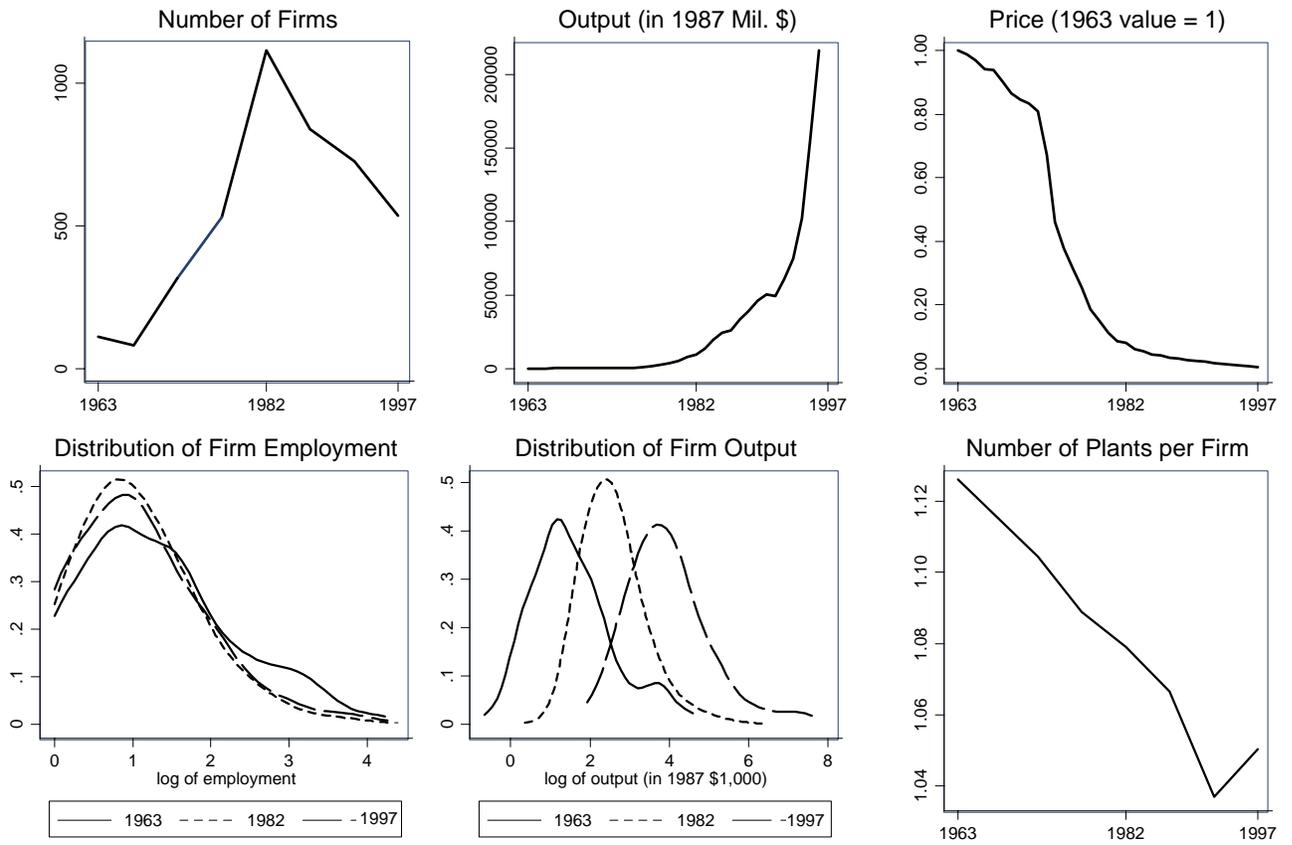
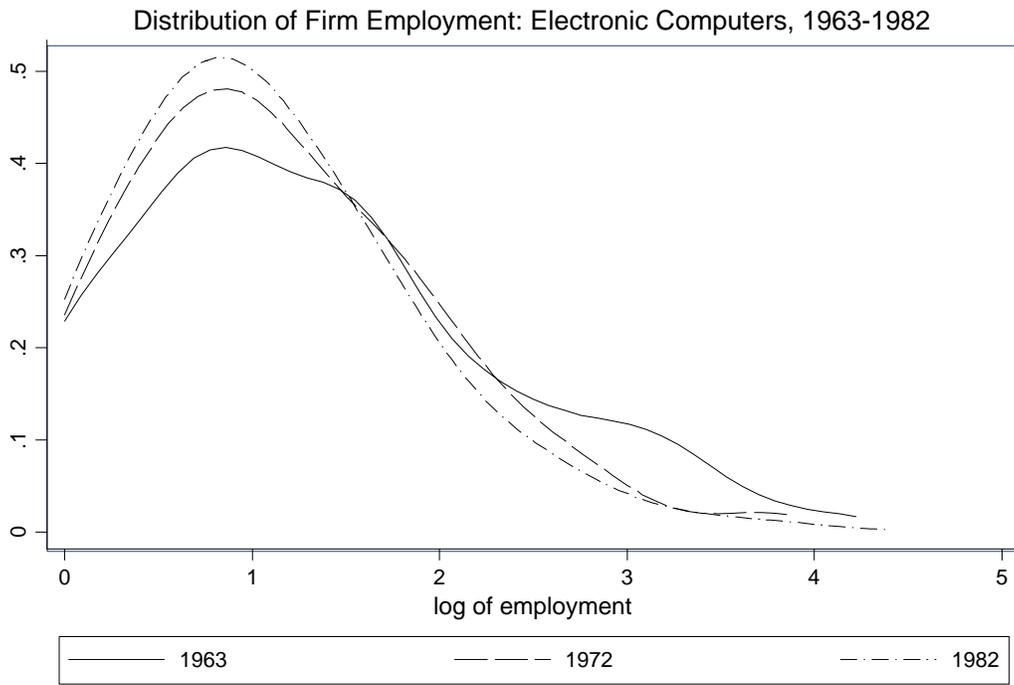
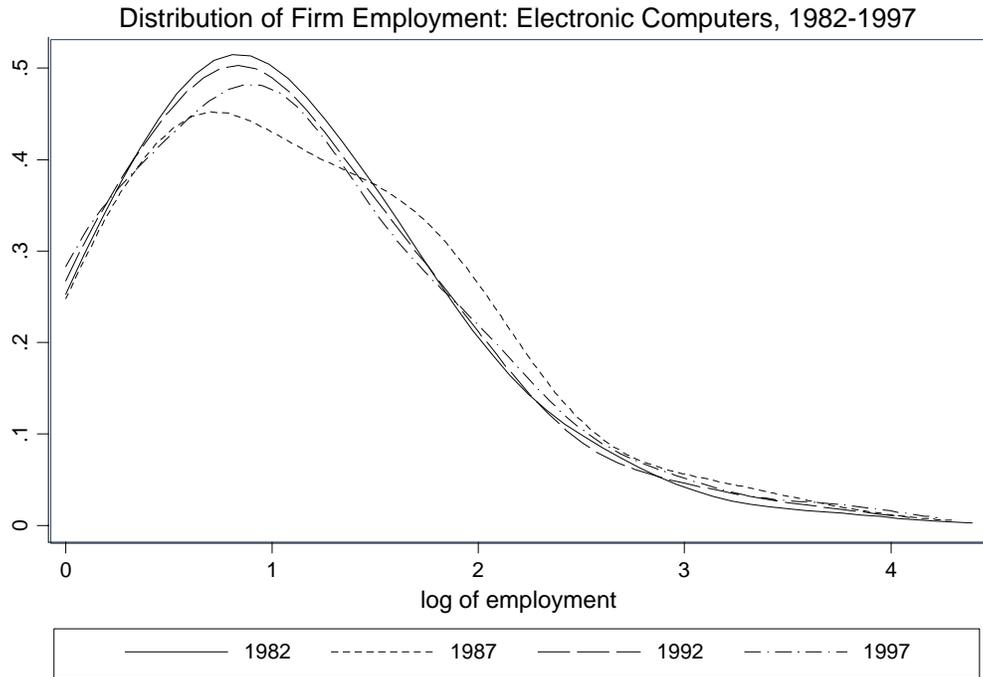


Figure 7: Evolution of key industry aggregates: Electronic Computers, 1963-1997.



Notes: Density estimates use a Gaussian kernel with a half-bandwidth of 0.3.



Notes: Density estimates use a Gaussian kernel with a half-bandwidth of 0.3.

Figure 8: Evolution of firm employment distribution: Electronic Computers, 1963-1997.

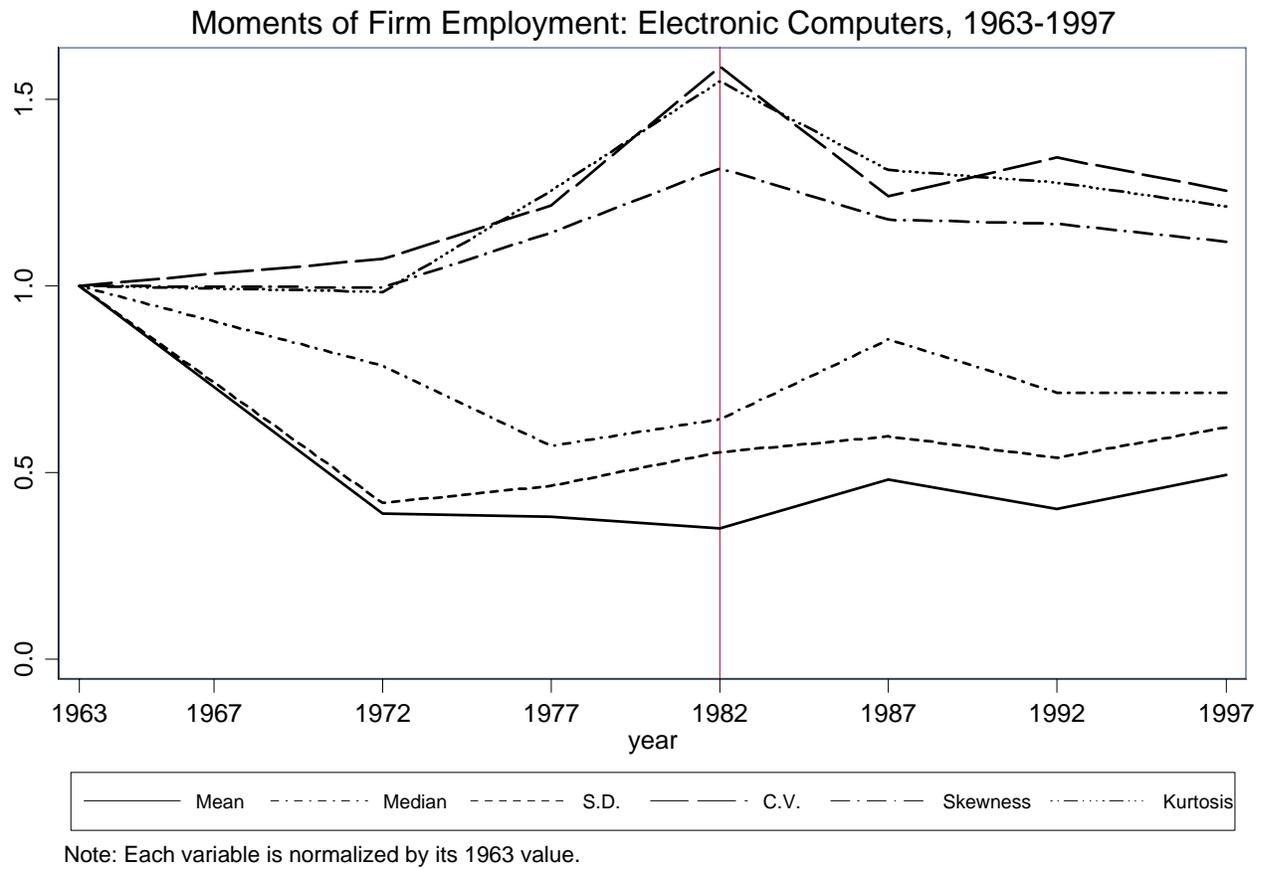
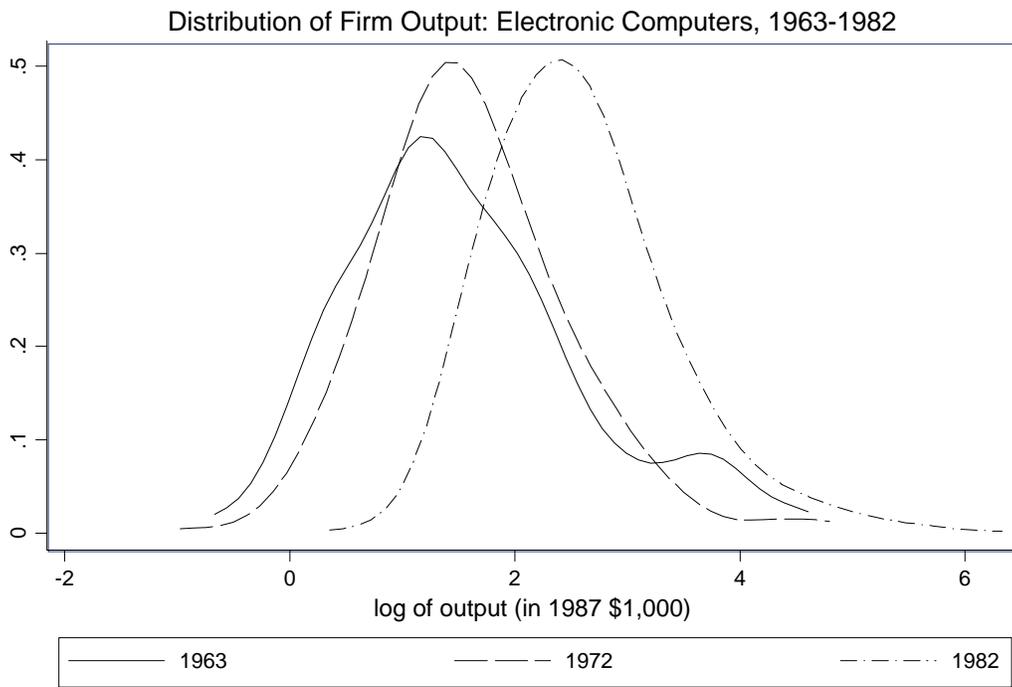
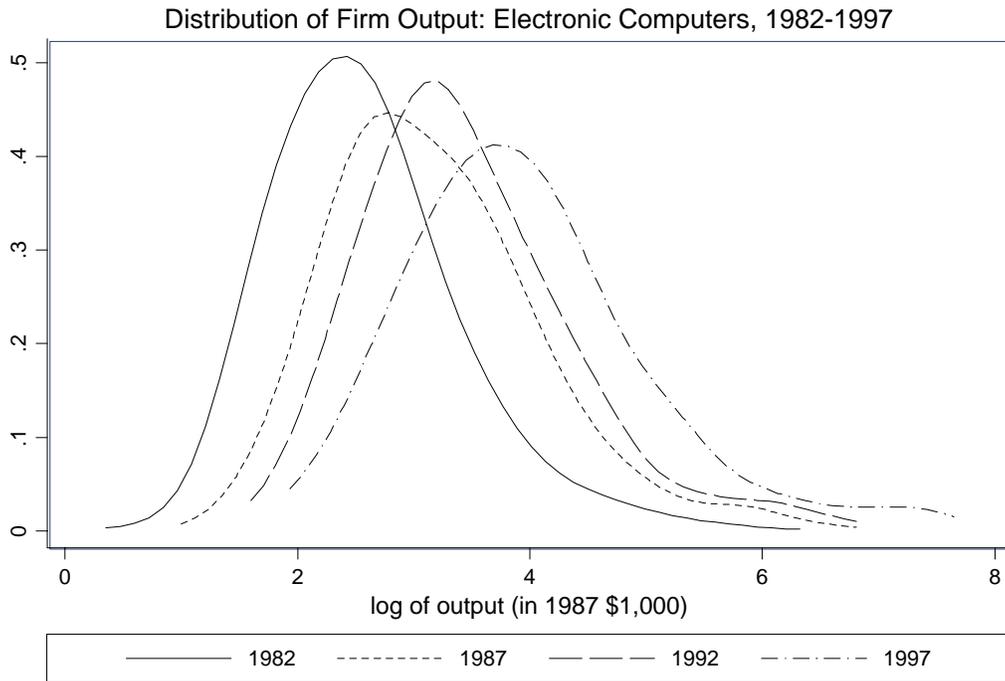


Figure 9: Evolution of the moments of firm employment: Electronic Computers, 1963-1997.

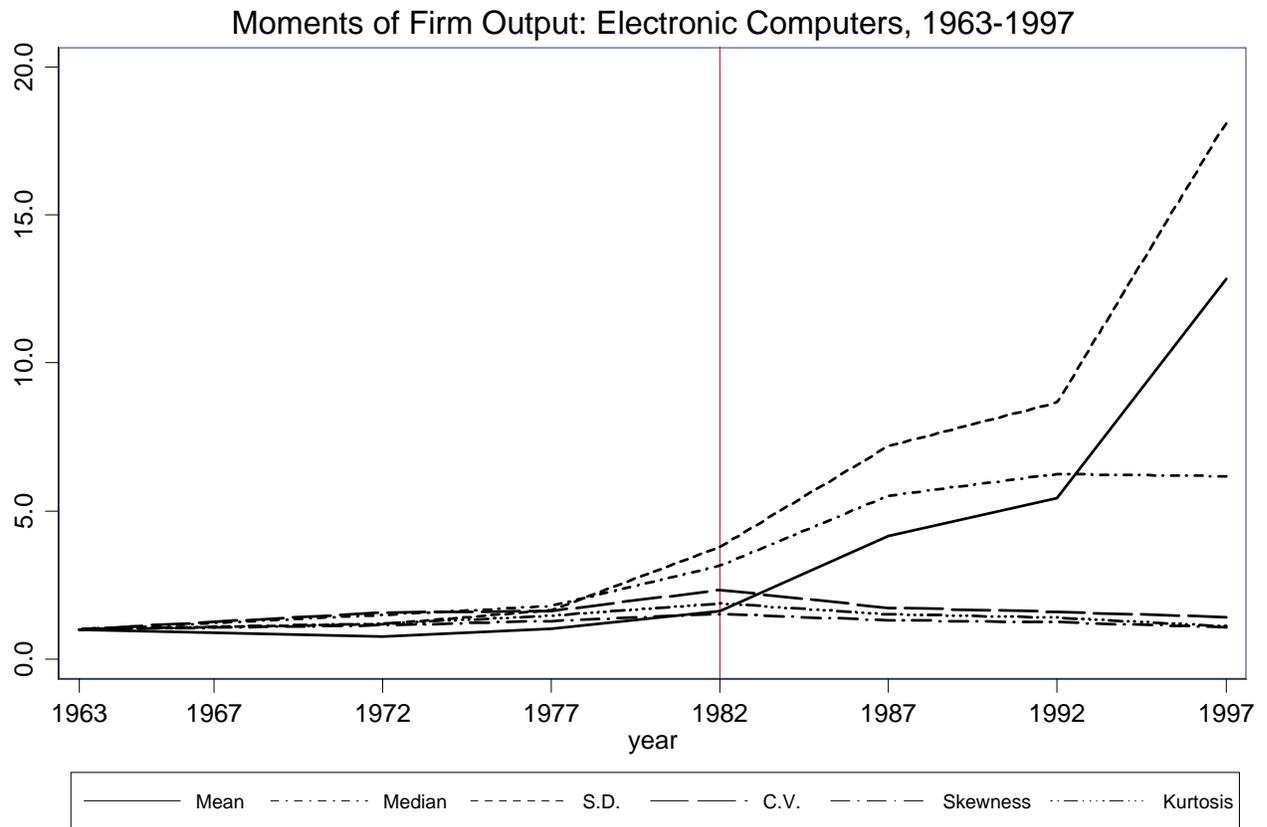


Notes: Density estimates use a Gaussian kernel with a half-bandwidth of 0.3.



Notes: Density estimates use a Gaussian kernel with a half-bandwidth of 0.3.

Figure 10: Evolution of firm output distribution: Electronic Computers, 1963-1997.



Note: Each variable is normalized by its 1963 value.

Figure 11: Evolution of the moments of firm output: Electronic Computers, 1963-1997.

Evolution of Key Variables: SIC 3674 Semiconductors

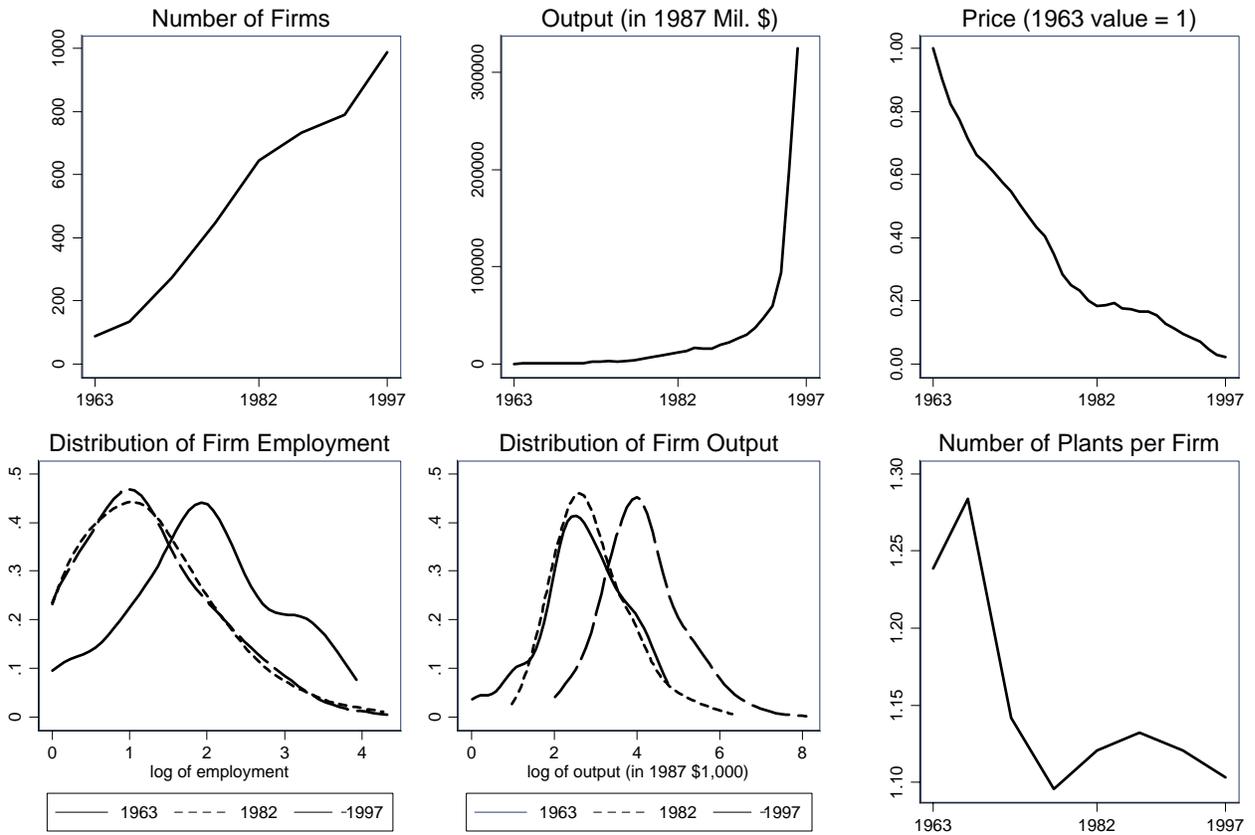


Figure 12: Evolution of key industry aggregates: Semiconductors, 1963-1997.

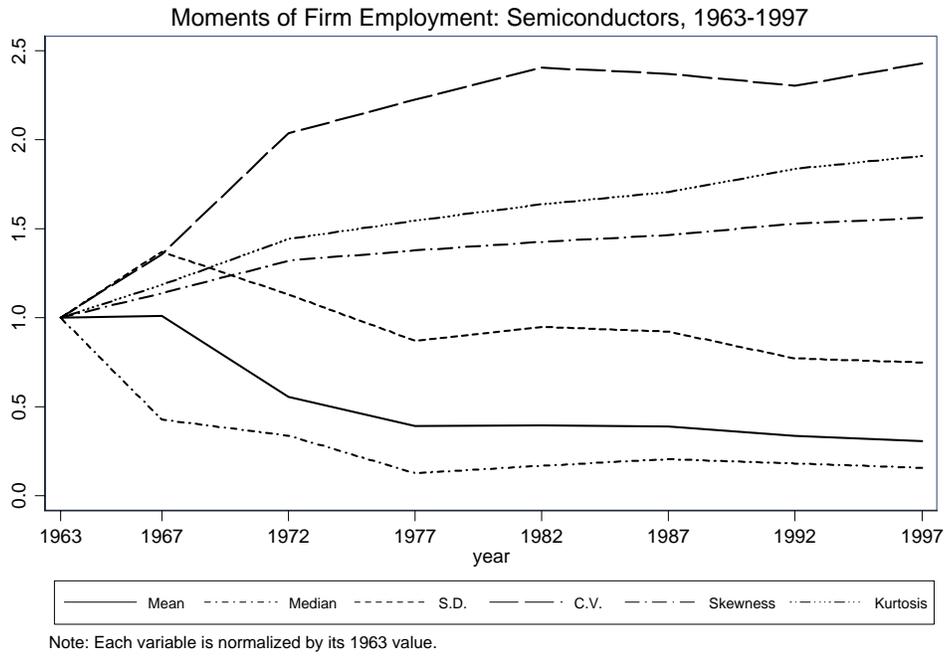


Figure 13: Evolution of the moments of firm employment: Semiconductors, 1963-1997.

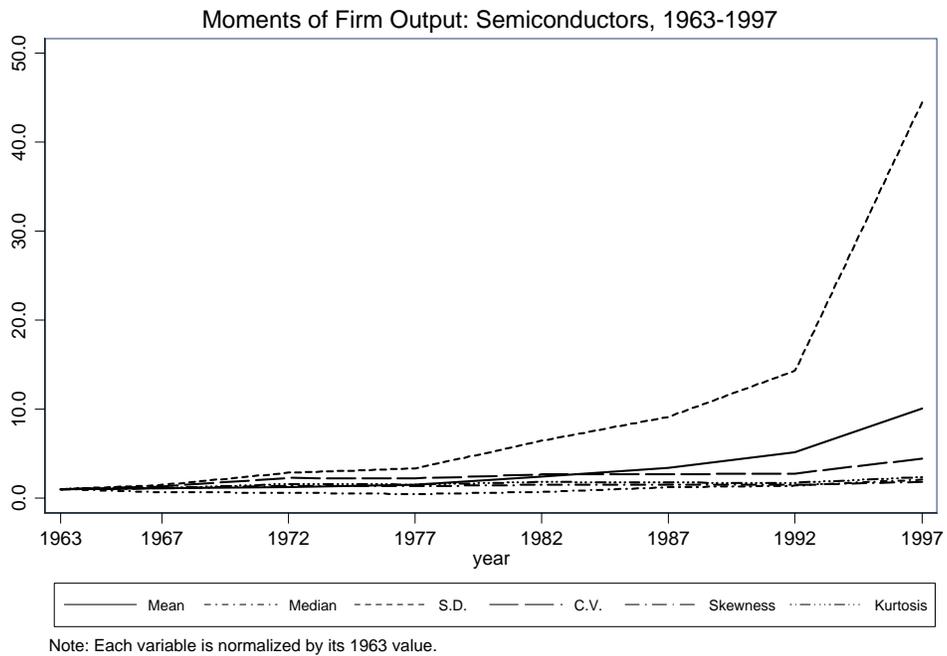


Figure 14: Evolution of the moments of firm output: Semiconductors, 1963-1997.

The evolution of the density of firm output

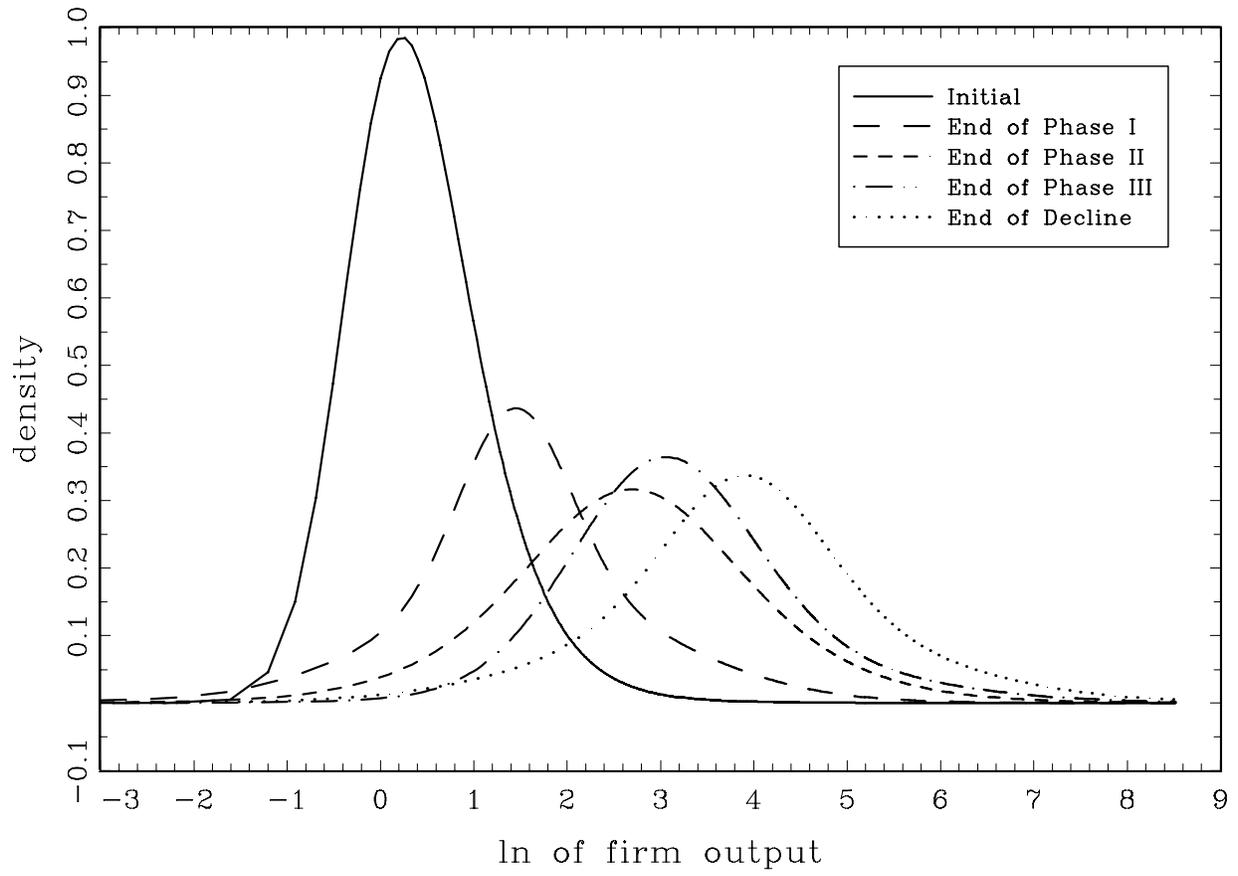


Figure 15: Representative evolution of firm output density along the life-cycle of an industry

The evolution of the density of firm employment

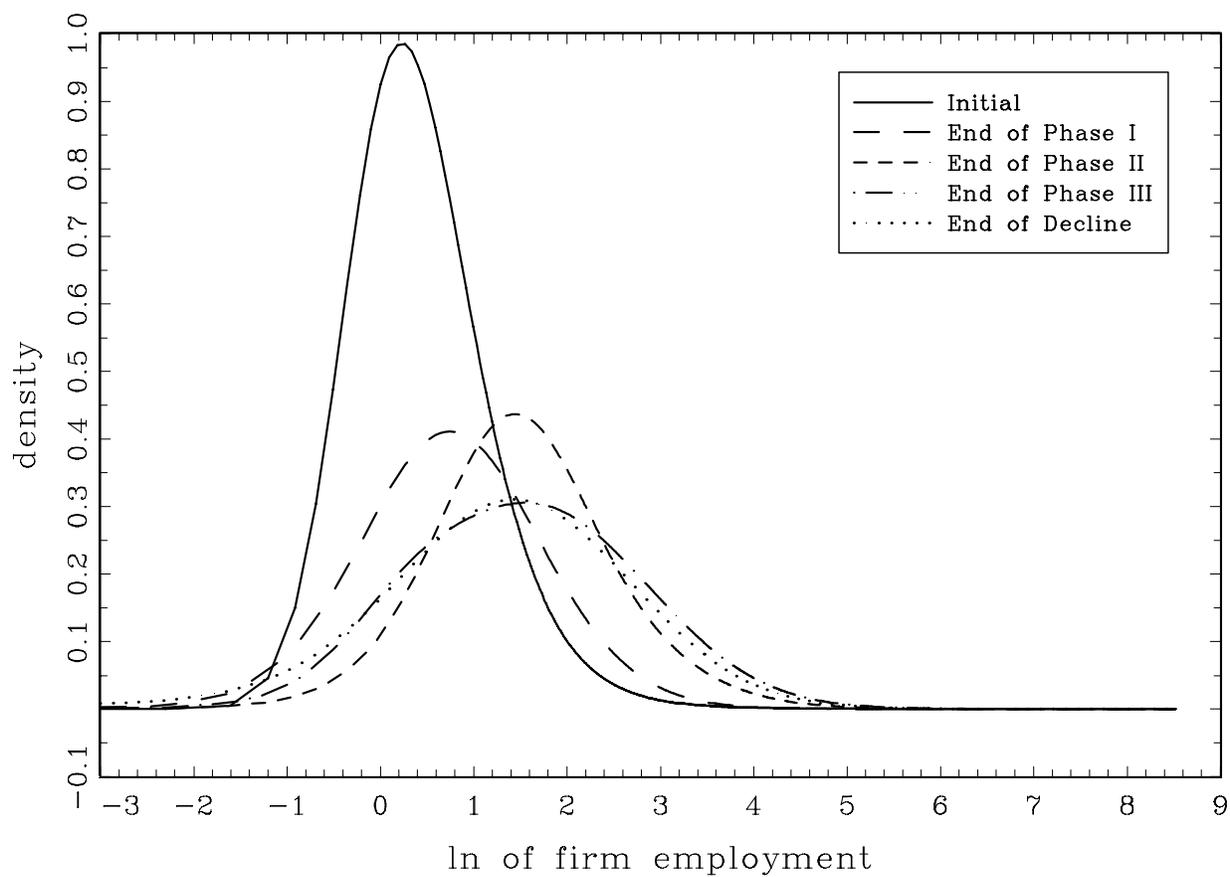


Figure 16: Representative evolution of firm employment density along the life-cycle of an industry