

NBER WORKING PAPER SERIES

CREATIVE DESTRUCTION OR JUST PLAIN DESTRUCTION?:
THE U.S. TEXTILE AND APPAREL INDUSTRIES SINCE 1972

Jim Levinsohn
Wendy Petropoulos

Working Paper 8348
<http://www.nber.org/papers/w8348>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 2001

The research in this paper was conducted while the authors were research associates at the Carnegie Mellon University Census Research Data Center. Research results and conclusions expressed are those of the authors and do not represent the opinions or official findings of the U.S. Bureau of the Census or the Carnegie Mellon University Census Research Data Center. The authors thank the Russell Sage Foundation for generous research support. We are grateful to many colleagues for helpful comments, but offer special thanks to Gordon Hanson and Amil Petrin. The views expressed herein are those of the authors and not necessarily those of the National Bureau of Economic Research.

© 2001 by Jim Levinsohn and Wendy Petropoulos. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Creative Destruction or Just Plain Destruction?:
The U.S. Textile and Apparel Industries since 1972
Jim Levinsohn and Wendy Petropoulos
NBER Working Paper No. 8348
June 2001

ABSTRACT

Are the U.S. textile and apparel industries examples of creative destruction or are they just plain destructing? We investigate this question using both aggregate industry-level data and plant-level data from the U.S. Census' LRD. We find that while the aggregate-level evidence is consistent with the common view of these industries as examples of declining industries, the plant-level data support a very different and much more hopeful view. We find that in the face of intensified international competition, each industry has evolved in its own way. In textiles, there has been tremendous capitalization. In apparel, the organization of production has changed. In both cases, industry productivity has increased markedly, and this is mostly because individual plants are becoming more productive.

Jim Levinsohn
Ford School of Public Policy
University of Michigan
Ann Arbor, MI 48104
jamesL@umich.edu

Wendy Petropoulos
Lexecon
332 South Michigan Ave., Suite 1300
Chicago, IL 60604
Wpetropoulos@lexecon.com

1 Introduction

Are the U.S. textile and apparel industries examples of creative destruction or are they just plain destructing?

The answer matters. It matters for the livelihoods of the over 1.5 million Americans who still work in these industries. It matters for the policy makers who negotiate the trade policy that sets the rules for textiles and apparel trade— trade that comprises about ten percent of all world trade in manufactures. And it matters to those who hold these industries up as examples of the changing nature of production in U.S. manufacturing, for the answer speaks to the impact of outsourcing (in the case of apparel) and capitalization (in the case of textiles) on industry outcomes.

Answering this question means defining creative destruction and plain destruction. We have a pretty good idea what the latter means on several dimensions of industrial change. It means plants close, employment shrinks, output declines, and productivity stagnates. Policy toward declining industries is usually focused on easing the pains of adjustment. Smart public policy for a declining industry usually does not advocate promoting entry.

Defining creative destruction is trickier. The phrase dates back to Joseph Schumpeter (Schumpeter (1942)). He wrote of an economic structure, “incessantly destroying the old one, incessantly creating a new one. The process of Creative Destruction is the essential fact about capitalism.” But just what does this mean for an industry’s plants, employment, output and productivity? The answer is not spelled out.¹ In terms of jobs, we take creative destruction to mean that although many jobs may be disappearing, new and hopefully better jobs are simultaneously replacing some of those lost. In terms of establishments, it means that even as exit occurs, there is substantial simultaneous entry. In terms of productivity, it means that the least efficient firms exit while survivors become more productive and clever new firms enter the market. Public policy can be crafted to facilitate or at least not hamper creative destruction. Policy prescriptions are about opposite those for declining industries. Rather than easing adjustment, policy may be directed at enhancing adjustment.

The aggregate-level evidence suggests these industries are just plain destructing. Thousands of plants have closed. Hundreds of thousands of jobs have been lost. Real wages have declined. Import competition has stiffened. By many metrics, the U.S. textile and apparel industries appear to be classic examples of declining industries. In an evolving international economy, comparative advantage is a dynamic notion, and it may be that the U.S. can no longer compete in these industries.

¹In this paper, we consider the process of creative destruction *within* the textile and apparel industries, although we recognize that the notion is often applied in a wider context. i.e. We’ll look at plant-level outcomes within an industry while some might instead examine industry-level outcomes within an entire economy.

A more hopeful interpretation of the aggregate evidence is that the decline shown in the aggregate statistics is just part of the process of creative destruction.² Investigating whether there is in fact creative destruction occurring behind the rather sobering aggregate statistics requires dis-aggregated data.

In this paper we use plant-level data from the U.S. Census over a 20-25 year span to investigate these competing views of the textile and apparel industries.

We consider these two industries in tandem for several reasons. Public policy often encompasses both. For the last 25 years, the arguably single largest piece of international trade protection, the Multi-Fiber Arrangement (MFA) was focused on these industries. Legislated in 1973 on the heels of the Short-term Agreement on Cotton and Textiles (STA) and the subsequent Long-term Agreement on Cotton and Textiles (LTA), the MFA was set to be the *really* long term agreement governing global trade in textiles. In the 1993 Uruguay Round of the GATT, though, it was agreed that in 1995 the MFA would begin a ten year phase-out. While the MFA is sufficiently complicated that quantifying exactly its impact on the U.S. textile and apparel industries is impossible, the past 30 years have certainly witnessed substantial protection in these industries, and only in the last 5 years (years that post-date our sample) have trade restraints begun to be systematically lowered.

The MFA is not the only reason these two industries are so frequently linked. The largest consumer of textile industry products made in the U.S. is the U.S. apparel industry, and textile industry products are the largest input to the apparel manufacturing process. Because of this relationship, the competitive forces facing either industry, and the success or failure of each industry's response, affects the other.

One such force with which both industries contend is the steady increase of globalization. Since the early 1980's, the developing and newly industrialized economies (NIC's) have become economically larger and more open. With this increased openness has come intensified import competition for U.S. firms. Competition from newly industrializing countries is particularly fierce for textiles and apparel; these industries have historically been driving forces for economies undergoing industrialization because of their labor intensiveness and relatively simple technologies that can be operated efficiently at small scale. Increases in openness from developing and NIC's has also brought enhanced outsourcing possibilities for U.S. manufacturers. The global economy

²Schumpeter himself warned against drawing conclusions about an industry based on a snapshot of aggregate data. Using language more florid than that permitted economists today, he warned that:

[S]ince we are dealing with an organic process, analysis of what happens in any part of it— say in an individual concern or industry— may indeed clarify details of mechanism but is inconclusive beyond that. Every piece of business strategy acquires its true significance only against the background of that process and within the situation created by it. It must be seen in its role in the perennial gale of creative destruction; it cannot be understood irrespective of it or, in fact, on the hypothesis that there is a perennial lull.

simply figures into the business plans of firms in both of these industries today in a way it did not in the early 1970's.

These two industries have also faced changes in the way their products are manufactured. In the apparel industry, there has been a shift toward a quick-response production structure, so that a piece of apparel can go from design to cutting to sewing to delivery in a fraction of the time such processes used to take, minimizing the inventories of work-in-progress that plants have on hand. But there has not been extensive mechanization of the production process in apparel; while the cutting of fabric has been mostly computer controlled since the early 1980's, apparel production remains relatively labor-intensive as the flimsy nature of fabric makes automation in this industry trickier than in other industries. Probably the biggest change in the way that apparel is manufactured relates to the tremendous growth of outsourcing, especially during the 1990's. This increase in apparel outsourcing also directly impacts the textile industry since much of the outsourced apparel production still relies on (now exported) textiles from U.S.-based mills.³

In the textile industry, the biggest change in the way products are manufactured relates more directly to a shift in technology. In almost all stages of textile production, from yarn spinning and slashing to fabric formation and finishing, new technologies have become available and have been widely adopted by U.S. firms. In 1976, there were about 275,000 shuttle looms in the U.S. textile industry. By 1997, there were about 10,000. In their place, large, highly capital-intensive plants using shuttle-less looms have become the norm. Other technical changes within the plant include wider looms, the increased use of robotics to move material within the plant, and substantially higher extrusion rates for man-made fibers. Another change in textiles involves the introduction of new products. Polartec and Gore Tex are two examples of consumer textiles that did not exist in 1970. There are many other new high-tech textiles for industrial use. Thanks partly to this increased mechanization of production and the introduction of new products, the U.S. textile industry has been successful at exporting a substantial and increasing amount of output over the last two decades, even as it has faced increasing imports.

The broad picture is one in which two industries shared a common fate in terms of the impacts of globalization but each industry responded in its own way. The apparel industry has moved much of the lower-skilled portions of the industry outside of the U.S. The textile industry has responded by substantial capitalization and investment in newer technology. In light of these varied responses, we investigate whether these industries have undergone a process of creative destruction or whether they have simply adjusted to their fates as declining industries.

We do not preclude the possibility that despite the common threads binding these two industries, one may decline while the other may re-invent itself. In fact, though, we find little

³Between 1995 and 1999, the value of textile and apparel imports falling under production-sharing provisions of U.S. trade law increased by 60%. This increase doesn't include the growth in outsourcing to NAFTA or CBI nations, which has also been substantial.

evidence of this sort of a split. Rather, we find that while the aggregate data in both industries look bleak, the plant-level data paint a much more Bruegel-like picture. There is a lot going on in these industries and while making sense of it all is challenging, careful analysis of the data suggests that the U.S. textile and apparel industries provide pretty good examples of creative destruction. In the face of intensified international competition, these industries have tremendous entry, much job creation, and surviving firms are becoming much more productive.

The organization of this paper is as follows. In the next section, Section 2, we provide some background on the U.S. textile and apparel industries. This section discusses some of the recent transformations these industries have undergone, what sort of products are produced, and some of the most salient changes in policy and technology. In section 3, we introduce the data on which we rely– the Census’ Longitudinal Research Database. In Section 4, we examine implications of creative destruction for plants, jobs, and productivity. Section 5 concludes.

2 The U.S. Textile and Apparel Industries: Some Background

The purpose of this section is to provide background on some of the more economically salient changes the textile and apparel industries have undergone in the last 25 years.⁴

U.S. industrial statistics treat the textile and apparel industries as separate “2-digit” industries– SIC 22 and 23 respectively. Each of these industries contain a diverse set of goods, goods that are produced by a wide variety of production processes and are sold to a broad array of final demanders. Plants falling within the textiles industry include those that bleach, dye or texturize fabrics for other plants on a commission basis; those that weave thread or yarn into fabrics sold to apparel manufacturers or makers of industrial machinery; and those that knit yarns into outer-wear sold directly to retailers. The traditional process for fabric formation includes three steps, which are often performed by different plants. The first step is picking fibers from bales to attain desired blends, and spinning the fibers into thread or yarn. The second step is weaving or knitting the thread or yarn into fabric, and the final step is chemical processing to finish the fabric. More advanced, non-woven fabrics are produced instead by mechanically, thermally or chemically bonding or interlocking fibers, filaments or yarns. In spite of the variety of products and production processes included in this category, the textiles sub-industries have undergone some of the same broad changes, including technology growth embodied in new machinery or new products.

The over-arching change in the textiles industry has been the move toward greater capitalization. Production has become a very capital-intensive process and this has induced significant

⁴The interested reader is referred to Abernathy, Dunlop, Hammond, and Weil (1999)– a very insightful overview of the recent transformations that have taken place in these industries. Another careful and detailed, but slightly more dated, study of these industries is William Cline’s “The Future of World Trade in Textiles and Apparel.” (Cline (1989)). Cline’s study is much more focused on the role of international competition.

recent investment. Indeed, from 1972 to 1992, the capital stock per employee in textiles has risen over 50 percent in real terms.

The apparel industry contains plants producing an equally wide variety of goods with a plethora of production processes. In addition to clothing items like shirts and trousers, this category includes apparel belts, embroideries, house-furnishings and automotive trimmings. The traditional apparel production process involves first designing the product and choosing fabrics, then pattern making, grading, cutting, sewing and pressing. Purchasers of apparel goods include retailers, but increasingly, apparel firms are retailing their own goods through outlet chains and catalogs.

In the apparel industry, production-related changes have arisen less from technical advance than from changes in the optimal organization of production. While there have been important technical advances in this industry, including the computerization of tasks such as product design and the placement of patterns on fabric to minimize waste, technical changes have been less important in apparel than in the textiles industry.

Over our sample period, however, there have been extensive changes in the organization of the apparel production process. One change, which we shall examine further in the following section, has been the increase in outsourcing abroad by U.S. manufacturers. Special exemptions available for goods that have been assembled abroad from U.S.-made components have been encouraging producers to move the assembly part of the production process to countries with cheaper labor over the last three decades.

In the last quarter of our sample period, another important change in apparel production has been driven by the demands of apparel retailers moving toward “lean retailing” practices.⁵ “Lean retailing,” spearheaded by large retailers, involves decreasing retail inventories by keeping careful (computerized) track of inventory levels and sales of every single item carried, and restocking specific items as needed by size and color. To make lean retailing work, however, retailers require that manufacturers be able to respond to their orders for re-stocks within days or weeks, rather than months or even once per season as had been the norm. The pressure for change these demands put on apparel manufacturers is hard to overstate. Manufacturers could accommodate the retailer demands for quick response times by either drastically increasing their own inventories, a costly move in terms of the inventory risk involved,⁶ or by changing their production process to speed the production of a complete apparel item. The production processes widely used by apparel manufacturers at the beginning of our time period, the progressive bundle and straight-line systems, involves breaking up each item’s assembly into a large number of highly

⁵These practices, and their effects on manufacturers described below, are also important forces in some of the textiles sub-industries that produce goods that go directly to retailers, for example, some knits as well as domestic goods such as sheets and pillow-cases and towels.

⁶Lean retailing raised the risk associated with even the initial level of inventories, since retailers would re-order only the specific colors and sizes of goods in demand, rather than the whole line of a product. This is discussed in Abernathy, Dunlop, Hammond, and Weil (1999).

specialized tasks. While the rate of output for these systems is high, it can take a long time for an individual item to pass through all stages of the production process, making it difficult for firms to respond quickly to orders for specific goods. Many manufacturers found it necessary to switch to a modular production system, breaking up production into fewer steps and requiring each worker to perform a wider variety of tasks. Modular production allows manufacturers to respond to retailer demands for quick response times without assuming unbearable amounts of inventory risk.⁷

As noted in the introduction, international competition now figures into the business plans of firms in these industries in a way that it did not 30 years ago. In an attempt to quantify some aspects of globalization in these industries, Table 1 presents some data on U.S. total imports of textiles and apparel as well as a measure of outsourcing. These are *quantity*-based measures of import competition and export performance.⁸ In textiles, the real value of imports was relatively constant from 1972-1982. Beginning around 1983, imports began a steady increase, and they have tripled in real terms from 1982-1994.⁹ The increase in U.S. imports came from both developed and developing countries with only a slightly greater amount of the increase coming from the latter. The U.S. textile industry has, throughout this period, exported to the rest of the world. In 1980, exports almost equaled imports. Since then, while exports have continued to increase, imports have grown faster, and the U.S. now imports about twice as much textiles as it exports.

In apparel, the picture is quite different. The absolute volume of exports is surprisingly high—apparel and textile exports are of comparable value throughout much of this period. Imports, on the other hand, have exploded, increasing almost ten-fold in real terms from 1972 to 1997. Whereas the increase in textile imports came from developed and developing countries, the increase in apparel imports is mostly from the developing countries. The import and export figures in Table 1 are in terms of real values. If instead one measures imports by physical

⁷For a much more detailed description of these changes see Abernathy, Dunlop, Hammond, and Weil (1999).

⁸The data in Table 1 are compiled from multiple sources. The import and export data in the literature are remarkably varied and seems to depend on which SIC revision is used, which, if any, deflators are used, and whether one uses USITC data among other things. Because of this variation, we want to be clear on how the data in Table 1 were constructed. The imports and exports data are from the Feenstra database available on www.nber.org. We used the file `sic58_94.asc.zip`. This data, though, uses the 1972 SIC revision. In order to make the data comparable to our LRD data, the Feenstra data were translated into the 1987 SIC revision using the concordances available with the NBER-CES Manufacturing Industry Database available at www.nber.org. Also, the Feenstra data stops in 1994. See (Feenstra (1996)) for details. Data for 1995-97 are from the U.S. Department of Commerce, International Trade Administration. Finally, the data were converted from nominal to real figures using the SIC-specific price deflators available on the NBER-CES Manufacturing Industry Database. A Stata program to create the import and export data in Table 1 is available on request. The data for Item 807 imports for 1972-1980 are from Research Department, International Ladies Garment Workers Union, and those for 1987-1995 are from the U.S. Department of Commerce.

⁹The data for 1995 onward are from a different source than that for 1972-1994 and this accounts for the downward blip in apparel and textile imports in 1995.

quantities (typically in million square meter equivalents in these industries), the message is the same.(OTA (1999)) Import competition, as measured by either value or physical units, has increased substantially with the apparel industry the harder hit.

There are strong arguments that one should measure import competition by prices, not quantities. Analysis of the Bureau of Labor Statistics import price indices for textiles and apparel shows that from 1982 to 1992, the prices of apparel rose only about half as quickly as those of textiles, suggesting again that import competition (judged by prices this time) was more intense in apparel than in textiles.

While total imports of textiles and apparel increased from 1987 to 1997, the nature of these imports changed dramatically over this brief period. As indicated in Table 1, about 5 percent of apparel and textile imports were what is referred to as Section 807 imports in 1987. These are goods that are essentially re-imported after their components were first exported and assembled abroad, and duty is only paid on the value-added. The vast majority of this trade is with either Mexico or the Caribbean Basin countries. These imports are one (imperfect) measure of outsourcing.¹⁰ In only 10 years, Section 807 imports (now covered under tariff provision 9802.00)¹¹ of U.S. manufactured goods that were assembled abroad increased over five-fold and these imports now account for about a fifth of all textile and apparel imports. Section 807 trade with Mexico rose especially quickly after the passage of the NAFTA.

This trade (now renamed) is expected to increase even more with the May 2000 passage of the Caribbean Basic Economic Recovery Act (CBERA). This act, which began to phase into effect in October 2000 essentially extends NAFTA-like treatment to Caribbean Basin nations in terms of textile and apparel trade. In sum, outsourcing abroad (by this measure) appears to be large, increasing, and fairly recent.

In light of the intensified international competition as well as the changes in technology (in the case of textiles) and industry organization (in the case of apparel), we seek to examine how plants have responded. As posed at the outset, have these industries become declining industries

¹⁰One might think that there would be well-established data on outsourcing. There is not. As noted by the U.S. International Trade Commission:

Official U.S. statistics are becoming increasingly unable to quantify the magnitude and scope of production-sharing activity since a significant and growing portion of imports from production-sharing operations does not enter under Chapter 98 provisions because goods are eligible for duty-free treatment under other agreements or tariff-preference programs.

(Rodriguez-Archila (2000)). This is particularly true in textiles and apparel where much of the outsourcing involves either NAFTA or Caribbean Basin countries.

¹¹The domestic content of U.S. imports entered under this provision is exempt from both duties and the Customs merchandise processing fee (the so-called user fee). Articles that are imported free of duty, either under a Most-Favored-Nation agreement or under a variety of trade preference programs, have a greatly reduced incentive to enter under the 9802/807 production-sharing provisions. While imports of duty-free articles continue to use provision 9802 because of exemption from the user fee, the value of 9802 imports may understate the amount of imports using U.S. components in their production.

or are they examples of the sort of creative destruction Schumpeter imagined? To address this issue, it is necessary to turn to plant-level data. We next describe our data sources.

3 Data: The LRD

The data we use come from the U.S. Census' Longitudinal Research Database (LRD), a panel data set constructed from reports that the Census collects directly from U.S. manufacturing establishments.

The data set is made up of two different surveys. The Census of Manufacturers is conducted every 5 years, and includes all U.S. manufacturing establishments with some minimal number of employees. These establishments are sent a short form survey requesting basic information about their operations and organization for the preceding year. For plants with fewer than the minimum number of employees, which are not required to file a report, the Census Bureau uses instead Social Security and other Internal Revenue Service records to construct some of the variables for these plants, called "administrative records", and imputes the rest. The second survey is the Annual Survey of Manufacturers (ASM), which covers a sub-sample of plants in operation and is conducted annually. For this sample, plants with more than 250 employees, and those making up a significant portion of the shipments in their industries, are selected with certainty, while smaller plants are selected with some probability. The ASM sample is designed to cover as much of the productive activity in the industries as possible, not to present a representative sample of plants, and the selection criteria for the samples have varied across industries and over time. Additionally, every five years a new ASM sample is chosen, and to minimize the reporting burden on small plants, those that were covered in one sample are less likely to be covered in the subsequent one.

Our panel spans the years 1972-1995, and contains five Census years (years ending in "2" or "7"). Most of our analysis focuses on the census-year data. The data contain both balance sheet (stock) and revenue statement (flow) variables, including information on levels of factors of production (such as employment payrolls, worker hours, and costs of fuels, electricity and materials) and on the products produced. In addition, unique establishment-specific identification numbers allow us to link plants across years. These identification numbers are invariant to changes in plant ownership; plants that are acquired by other firms, or plants whose form of organization changes retain the same identification number throughout.

4 What Actually Happened? Plants, Jobs, and Productivity in the U.S. Textile and Apparel Industries

In this section, we document the changes the textile and apparel industries have undergone. We examine what has actually happened to plant entry and exit, job creation and destruction, and productivity in the textile and apparel industries using establishment-level data. To the extent possible, our approach is to first provide the sort of publicly available aggregate evidence on the topic at hand, draw preliminary conclusions on whether this evidence is suggestive of a declining industry or one undergoing creative destruction, and then turn to the plant-level data. There, we investigate whether the plant-level data further substantiates the message of the aggregate data or whether they contribute to a new and different perspective.

4.1 Establishments

We begin with an examination of the number of plants in the textile and apparel industries. Declining industries, as well as those undergoing creative destruction, exhibit substantial exit. The difference between the two is that in industries undergoing creative destruction, the exit is, at least in part, countered with simultaneous entry.

Table 2 provides count data on the number of plants in the census years. The first column includes all plants in the data set. Plants that appear only as Administrative Records (ARs) are included here. (Recall, these plants are small plants that are not required by the Census to file a report. Rather, the Census gets some information from administrative data such as IRS data and imputes the rest.) We will typically exclude AR data from our subsequent analyses. One reason for this is that the unique code that identifies a plant over time, and hence allows us to create our panel, seems to disappear and re-appear fairly randomly from the AR data. This is not true entry and exit, but rather mostly an artifact of reporting errors. Including the ARs in our entry and exit figures would artificially inflate estimates of entry and exit.

More reliable data excludes the plants that appear only as administrative records, and these counts are given in the “No AR” columns.¹² Regardless of whether one includes or excludes the AR-only plants, the number of plants is declining in apparel and textiles.

In apparel, the decline is modest when tiny plants are included but quite substantial once we’ve deleted AR-only plants. This indicates much net exit by all but the smallest plants. Some of this exit represents plants that are going out of the business entirely while some of this exit represents plants that exit manufacturing but remain in the business as jobbers. The fact that

¹²We report the figures including the ARs in this table for two reasons. First aggregate published statistics often do not delete the AR-only firms and by providing AR counts, our data may be more compatible with other published sources. Second, by reporting both, readers can obtain an estimate of how many very small plants there are in the industries.

the smallest firms show less net exit also captures the notion that many of the entrants may be very small plants. With the less reliably reported data excluded, there is an approximately 25 percent decline in the number of apparel plants.

In textiles, there is a similar picture. Even including the AR-only plants, there is net exit on the order of 20 percent over the 20 year sample period. When the administrative records data are excluded, we find that there are only about two-thirds as many textile plants in 1992 as there were in 1972.

The fact that the relative declines are larger when only examining the non-AR plants is doubly sobering. Not only is the number of plants declining, but it is the larger more established plants that are showing greater net exit. While the pattern of net plant exit we find in the textiles and apparel industries doesn't hold across manufacturing overall (which includes such newer industries as computer chips, telecommunications equipment and semiconductors), it is echoed in many of the "older" manufacturing industries.¹³ We conclude that the industry-level evidence on plant counts is quite supportive of the view that these industries are simply declining industries.

To examine the case for creative destruction, one needs to look behind the net declines in Table 2 and examine the data on simultaneous entry and exit. Table 3 does this. For each of the five year intervals and using only non-AR data,¹⁴ we compute the gross rate of entry and of exit. The former is defined as the number of new plants divided by the number of initial plants in a 5 year interval while the latter is the number of exiting plants divided by the number of initial plants. The results are surprising and are especially so for apparel.

In textiles, a quarter of all plants were entrants between 1972 and 1977 while a third were exiters.¹⁵ Although these numbers fluctuate, they are broadly representative of the sort of gross changes in the number of plants in textiles. In apparel, the numbers are more striking. For example, between 1972 and 1977, the number of (non-AR) plants was virtually constant. Still, almost *half* of all plants in this period were entrants and almost half exited so that the net change was negligible.

It would be easy to examine the industry-level figures for the first 15 or so years of our sample and conclude that the effect of the MFA was to protect existing plants, and preclude any Schumpeterian creative destruction. It would also be dead wrong, although one needs plant-level data to ascertain this. In general, underlying the net exit of plants in these industries are

¹³Over the last 30 years, plant numbers have fallen in other manufacturing industries such as Food Products; Tobacco; Leather; Stone, Clay and Glass; and Primary Metals.

¹⁴Hereafter, all tables will only include non-AR data unless otherwise stated.

¹⁵We use the simple appearance or disappearance of a plant to indicate entry or exit. An alternative method of measuring entry and exit involves the use of "coverage codes." Because these codes are not especially reliable for our industries toward the end of our sample, we do not use them. Had we used coverage codes, we almost surely would obtain lower entry and exit rates, hence our estimates can be considered upper bounds on entry and exit rates.

large and somewhat steady rates of both entry and exit. Relative to manufacturing as a whole, these industries probably exhibit much greater than average entry (and exit) rates.¹⁶ It is also apparent that the waves of particularly large net exit are explained more by a fall in gross entry than by an increase in the gross exit rate. That is, when the number of plants in these industries most declined, it was more due to a shortage of entrants than a surfeit of exits. Examining this pattern at the 3-digit SIC level, we find that the average variance in entry rates within SIC sub-industries exceeds the average variance in rates of exit.

A statistical implication of this table is that when we later try to estimate how entrants and exiters differ, we are dealing with large samples. A policy implication is that interventions to slow industry decline need to focus not just on supporting existing plants but also on encouraging entry. Indeed, even though these industries faced vastly intensified international competition, the exit rates at the start of the sample were almost unchanged two decades later! A final implication of this table concerns product variety. We do not observe the attributes of the products that plants actually produce. A common modeling assumption is that plants produce differentiated varieties. This assumption may be less objectionable in textiles and apparel than in some other industries. To the extent that the assumption captures some of reality, there are a lot of new varieties being introduced (and lost), and this too seems a part of creative destruction.

4.2 Jobs

One reason we care about the entry and exit of new plants is that these statistics fore-shadow job dynamics. When plants exit, they no longer have the potential for creating new jobs (and lose those existing at closure) while plant entry holds out hope of future job creation. The aggregate-level data on jobs are presented in Table 4. The picture is bleak. From 1972 to 1992, employment in apparel declined almost 30 percent while in textiles employment declined over 35 percent. The absolute number of lost jobs on net, 720,000, exceeds (to pick an example that should resonate with an academic audience) the total employment of faculty in 2 and 4 year colleges and universities in the U.S. By almost any reasonable metric, these are the sort of huge declines that befit a declining industry.

Examining the number of production and non-production workers at the average plant indicates that it is production workers who have borne the brunt of these job losses. In apparel, the mean number of non-production workers has been constant at 13. The mean number of production workers, on the other hand, has fallen from 77 to 62. In textiles, where the average plant size is larger, the pattern is the same. Mean non-production workers is unchanged over the sample at 23 while the average number of production workers has fallen from 165 to 139.

¹⁶The “probably” is due to the fact that we do not use coverage codes in computing entry and exit. Baldwin (1999) reports that the average entry rate over 5 year periods for the U.S. was on the order of 11.3 percent. Even accounting for differences in measurement criteria, the entry rates for textiles and apparel appear quite large relative to this average.

The reasons for the declines in production workers varies across the industries and are closely related to the figures for mean real capital stock per employee. In apparel, outsourcing probably explains the concurrent decline in production (and not non-production) workers and capital stock. Some production jobs are being outsourced, and this is accompanied by decreasing capital intensity. Management jobs are impacted less by outsourcing and this is consistent with the unchanging average non-production employment. In textiles, machines are replacing some production workers. Capital per employee rises from \$26,562 to \$37,222, while the number of production workers falls. In summary, though, the aggregate-level evidence on net job loss is clear. There were 720,000 fewer jobs in 1992 than there were in 1972.

As noted in the discussion of establishments, a difference between creative destruction and just plain destruction is that in both cases jobs decline, but with creative destruction, many new jobs are simultaneously being created. This presages the important difference between the net change in jobs and the gross change in jobs. If there is no concurrent job creation, the two are identical. On the other hand, if four jobs are created for each five lost, gross and net changes are very different. In order to quantify these notions, we adopt the by now standard definitions of Davis and Haltiwanger (1992).

We denote the employment at plant i in census year t as x_{it} . We then define the average employment at the plant as:

$$\bar{x}_{it} = \frac{1}{2}(x_{i,t} + x_{i,t-1}). \quad (1)$$

Summing across all plants gives the average industry employment, X_t . Hence,

$$X_t = \sum_i \bar{x}_{it}. \quad (2)$$

The growth rate of employment at a plant, g_{it} , is given by:

$$g_{it} = \frac{x_{it} - x_{i,t-1}}{\bar{x}_{it}} \quad (3)$$

This formulation has the nice property that it easily accommodates births ($g= 2$) and deaths ($g= -2$) of plants.

We begin to address the concurrent roles of job creation and destruction in Table 5. Entries in Table 5 give the average fraction of existing jobs that were created or destroyed between census years. That is, $1/T \sum_{t=1}^T \sum_i (x_{it} - x_{i,t-1})/X_t$ where there are $T=4$ census years.

In each five year period between censuses, a certain fraction of all the jobs existing in that five year period were created during the period. In textiles, 17, 13, 21, and 25 percent of the production worker jobs were created in each of the four five-year periods, so averaged over five year periods, 19 percent of all the existing production jobs were created in previous 5 years. This

is the first entry on the first row of Table 5. Computing the same figure for apparel, we find that 30 percent of apparel jobs were created in the typical 5-year period. For non-production workers, the figures are higher— on average 28 percent of jobs were created over the five-year intervals in textiles and 42 percent in apparel. The analogous figures for job loss are given in the row titled “Gross Job Loss.” For example, among production worker jobs in apparel, on average, 42 percent of all the jobs that existed during a five year interval were lost over the course of that interval. The row titled “Sum of Gross Job Flows” is an especially important one. Figures in that row give the percent of existing jobs that were either created or destroyed in the five year interval. The bottom row gives the net change. For example, for production worker jobs in textiles, between census years, on average 19 percent of all jobs were created and 30 percent were destroyed, so the average net change was an 11 percent decline.

In trying to understand the employment changes in these industries, Table 5 has several key messages. First, job churning, as measured by the sum of job creation and job destruction is substantial in both industries but larger for apparel. In textiles, in *each* five year period, half of all production jobs are either created or destroyed. In apparel, the figure is 73 percent. When we examine only the net job flows, textiles and apparel are the same. Second, somewhat surprisingly, job churning is greater for non-production workers in both industries. One might have expected that with the growth in outsourcing, non-production jobs would be relatively stable while production jobs would be the ones disappearing. This is not the case. Both job creation and job destruction rates are higher for non-production workers than for production workers in textiles and apparel. The fairly minimal average decline in the number of non-production jobs in these industries (8 and 7 percent) hides tremendous simultaneous creation and destruction. Third, entrants account for a disproportionate amount of new jobs while exiters do the same for lost jobs. This is especially apparent in apparel, where about three fourths of all job loss is accounted for by plant exit. Fourth, it is differences in entrants and exiters that account for the differences across industries. If we were to confine our analysis to continuing plants, average rates of gross job flows in the apparel and textile industries would look almost identical. Lastly, we note that one of the forces behind the coincident entry and exit of plants is the growth of certain sub-industries or product classes within textiles and apparel that is occurring even though the industry as a whole is shrinking. In textiles, the Carpets and Rugs, and Knitting Mills sub-industries were growing in terms of employment over the years 1972 - 1992.¹⁷ In apparel, non-clothing product classes such as gloves and mittens, hats and caps, and apparel belts were growing in terms of employment and plants. Even so, looking within each sub-industry we still find substantial rates of both gross job inflows and outflows.

(Davis, Haltiwanger, and Schuh (1996)) list annual rates of gross job growth and loss as well as gross and net job flows for the years 1973-88 for all U.S. two-digit manufacturing indus-

¹⁷It is interesting to note that these are two industries that sell output directly to final consumers and to other producers outside of the apparel industry.

tries.¹⁸ These figures are not directly comparable to the 5-year averages we report, but are still informative. Compared to U.S. manufacturing overall, the patterns found in textiles are fairly representative. The rate of job growth and loss in textiles is slightly below the average for U.S. manufacturing but is not an obvious outlier. Apparel, on the other hand, is an outlier. Of the 20 2-digit industries, apparel has one of the three largest rates of job growth and loss as well as gross and net job flows.

The broader question being addressed is whether job patterns befit an industry undergoing creative destruction. We believe the answer is “yes.” For each industry and for each type of worker, gross job flows swamp the net job flows. Depending on industry and worker type, the sum of gross job flows is from 5 to 13 times as large as the net decline, and this is due to the surprisingly strong job creation taking place—job creation that does not seem to fit the usual definition of a declining industry.

An important part of the process of creative destruction in these industries concerns how production is being re-organized, and this reorganization also has implications for job (as well as capital) flows. Toward this end, we have also examined job growth rates and investment rates for plants that we classify as mainly “outsourcers” or “insourcers” compared to plants that don’t fall into these categories. “Outsourcers” are (approximately) the top quintile of plants in terms of the value of work they contract out as a percent of their total sales (about 2% of sales for textiles and 5% for apparel.) “Insourcers” are plants whose main product class (a 5-digit SIC code) is some variety of contract work. The results, presented in Table 6, are striking.

Consider first the textiles industry. Plants that are classified as either outsourcers or insourcers are “down-sizing” at a much more rapid pace than are other plants. The difference between outsourcers and other plants is especially stark in their growth rates of non-production worker employment and capital stocks; outsourcers shrink by about 25% over each five year interval in their employment of these inputs, while non-outsourcing plants only cut back their non-production worker employment by 7% and their capital stock by 10%. In the apparel industry, while both outsourcing and insourcing plants appear to be “down-sizing” more rapidly as well, the difference is much more pronounced between insourcing plants and others. Apparel industry insourcers are shrinking in terms of both labor and capital stock at a rate of about 25% every 5 years, while non-insourcing plants are shrinking by less than 6% in all categories of inputs. The reasons behind their striking rates of reduction in input use, however, likely differ between the outsourcing and insourcing plants. “Outsourcers” are finding it worthwhile to outsource more and more of their production abroad. Many of these plants are exiting manufacturing altogether, becoming wholly importers and contractors. “Insourcers”, on the other hand, are losing business to foreign firms that face lower labor costs and to assembly plants that U.S. firms are opening up abroad.

¹⁸See Table 3.1 in (Davis, Haltiwanger, and Schuh (1996)).

4.3 Productivity

As noted in the introduction, import competition has intensified in textiles and apparel. Technology, too, has changed. In this section, we investigate how productivity has evolved in response to these and other changes. One of the oft-cited benefits of international competition is that it forces domestic firms to become more productive.¹⁹ An oft-cited cost of international competition is that, in some cases, a country may lag behind the productivity of its competitors losing its competitive advantage over time, and with this comes lost jobs, closed plants, and declining wages.²⁰

The implications for the evolution of industry-level productivity in declining industries are straightforward in a closed economy. Declining industries exhibit productivity growth below the domestic average. In an open economy, the metric is global. Declining industries exhibit slower productivity growth than their international competition. Unfortunately, we do not have reliable measures of productivity growth in textiles and apparel around the globe. While we can, and will, measure productivity growth in the U.S., we are unable to evaluate whether these changes suggest a declining industry in a more global context. Rather, we simply examine whether the textile and apparel industries have lagged behind other U.S. manufacturing industries.

We are able, though, to investigate several aspects of creative destruction as evidenced from productivity dynamics. Creative destruction implies that: i) as plants exit, it is the least efficient that are most likely to go out of business; ii) surviving firms learn to become more productive; iii) the net effect of entry is to enhance industry productivity; and iv) that among surviving plants, more productive plants grow while their less productive competitors lose market share. In this section, we investigate all of these implications of creative destruction.

4.3.1 Measuring Productivity

Before presenting results, we must first address measurement issues. Unlike plant entry and exit, and employment, productivity is not directly observable. Examining productivity, then, means coming to terms with a way to measure it. In this sub-section, we discuss two ways to measure productivity— a traditional and easy-to-use way and a state-of-the-art way.

The traditional method of measuring productivity at the plant level is to compute value-added per worker. This method has at least five attributes to recommend it. First, it is exceedingly simple to calculate. Both value-added and the number of workers are reported in the plant-level data so measurement is a trivial exercise. Second, it is a measure that, by itself,

¹⁹We repeat this assertion without endorsing it. As micro-economists, it is not immediately obvious to us why this should be so in an industry with competitive domestic markets.

²⁰Some might view this as a benefit citing instead the resources freed up from the industries in which the country was no longer competitive.

means something. That is, it is not an index number but rather a dollar-based value that is readily interpreted. Third, this is how many government statistics measure productivity, so this measure makes comparisons of, say, the apparel and textiles industries to manufacturing at-large very straightforward. Fourth, this measure is not dependent on choices regarding functional form and econometric issues that are likely to have an arbitrary element to them. Fifth, it is a pretty intuitive measure. More output per worker seems like it ought to be a good thing if one doesn't think too hard about it.

The problems with value-added per worker as a measure of productivity arise if one does start to think a bit harder. In the presence of other inputs, focusing productivity measurement on just labor can be misleading. An extreme example illustrates this point. Suppose it was possible to purchase a machine that allowed a single person to run an entire textile plant, but that this machine cost several billion dollars. Output per worker would be huge, but as a business decision, purchasing the machine would be stupid. One needs to account for the other inputs into the production process and this quickly leads one to measures of total factor productivity. There is another, more subtle, problem with using output per worker to measure productivity. The use of output per worker as a productivity measure strongly biases one toward finding a trade-off between productivity changes and employment changes, and this fosters a problematic political dynamic. (This trade-off between "productivity" and employment, for example is highlighted in Cline's study of the textile and apparel industries.) Holding output constant, the only way plant-level productivity (by this measure) will increase is if workers are laid off. With more accurate measures of productivity, it is easier to have both jobs and increased productivity. The current state-of-the-art in measuring productivity was developed by Olley and Pakes (Olley and Pakes (1996)). These methods were then adapted and improved by Levinsohn and Petrin (Levinsohn and Petrin (2000)). We use the Levinsohn and Petrin (hereafter LP) methodology for estimating plant-level productivity. We proceed by giving a very brief description of this methodology (as well as some problems with simpler estimation methods) and refer the interested reader to the Appendix and to the original papers for more detail.

If one is going to estimate, rather than simply compute, total factor productivity, the simplest way to do so is to estimate a production function using Ordinary Least Squares and use the residual from such a regression as the measure of productivity. The problem with this oft-adopted approach was pointed out over 50 years ago in Marschak and Andrews (1944). In essence, Marschak and Andrews noted that plants that have a large positive productivity shock may respond by using more inputs. To the extent that this happens, Ordinary Least Squares estimates of production functions will yield biased estimates and, by implication, biased estimates of productivity. The usual "fix" to this econometric endogeneity is to use an Instrumental Variables estimator. This approach would require an instrument that was correlated with plant-level input choices and orthogonal to the productivity shock. With plant-level (as opposed to industry-level) data, there probably are no valid instruments. This very basic problem went

mostly unsolved until 1996 when, using data from the LRD and focusing on the telecommunications industry, Olley and Pakes made use of an imaginative blend of economic theory and semi-parametric econometric methods to solve the endogeneity problem. Levinsohn and Petrin followed with some simplifications and amendments that seem to allow the original Olley and Pakes insights to work better with a wider variety of data sets. A brief description of exactly how we estimate productivity is given in Appendix A. In that Appendix, we describe how we estimate productivity in the apparel and textiles industries. Readers wanting more detail should consult Levinsohn and Petrin (2000).

4.3.2 Aggregate Productivity

Table 7 provides the aggregate-level evidence on productivity. That table lists the Bureau of Labor Statistics (BLS) Total Factor Productivity (TFP) index for manufacturing overall²¹ as well as productivity indices specific to textiles and apparel. For each industry, we list the BLS TFP index as well as two indices that we compute. These are the LP econometrically estimated measure and, for the sake of comparison, a normalized measure of value-added per worker (in constant 1987 dollars.) For all indices, industry-level productivity is computed as an output share-weighted average of the plant-level productivities. We normalize each index to equal 1.00 in 1972 and compute each index for each census year.

There are four main messages in the industry-level productivity figures. The first message speaks to a comparison between the econometrically measured LP estimate and the value-added per worker measure. We find that the traditional measure of industry productivity, value-added per worker, gives an enormously misleading picture of productivity in these industries. By this traditional metric, productivity about doubled in textiles and more than doubled in apparel. The reasons for the large discrepancies between measures is different in the two industries. In textiles, the capital per worker increased about 50 percent over this period. Accounting for the fact that capital is not free will make the econometric estimate of productivity (which accounts for all inputs) smaller than the productivity measure based just on labor. This is what is largely behind the overstated productivity gains when one uses the output per worker measure. In apparel, capital per worker is falling, yet the measure of productivity that accounts for all inputs is much smaller than the value-added per worker measure. The reason has to do with the decreasing returns to scale in apparel. (In results not reported here, we find returns to scale are about .90 in apparel.) In apparel, plants are becoming smaller over time (both in terms of capital and in terms of labor.) With decreasing returns to scale and shrinking plant size, the average product of labor, which is what the value-added per worker measure captures, will increase.

²¹The BLS indices are from the BLS web site, www.bls.gov, and are taken from the table titled “Multifactor Productivity in U.S. Manufacturing and in 20 Manufacturing Industries, 1949-99.” The table is dated September, 2000.

Henceforth we focus on the LP estimates of productivity. The second message of the table is that even using a measure of productivity that accounts for all factors (either the TFP measure or our own estimates), productivity has risen in these industries. Based on our own estimates, productivity in textiles has risen substantially from 1972 to 1992— almost a 70 percent increase, while in apparel the increase is a more modest 26 percent increase. Given the technological advances in textiles and the relative lack thereof in apparel, the fact that productivity improvements were greater in textiles seems intuitive. The third message in the Table is that, with only one exception (1987 to 1992 in apparel), industry-level productivity never actually fell. Given the pervasive turnover and exit in these industries, this is perhaps surprising.

The fourth message from Table 7 concerns the comparison between the LP estimates and the BLS TFP measures. Both measures take into account all the factors of production. The LP estimates do not impose as much structure on the data in that they do not directly impose the underlying economic restrictions that justify the correctness of the TFP measure. Nonetheless, the two measures give broadly similar pictures of productivity, and we find this reassuring.

The aggregate-level data in Table 7 does not allow one to pass judgment on whether the industry-level evidence is consistent with the notion that the textile and apparel industries are declining industries. To begin to address this question, we need to compare productivity estimates for textiles and apparel with those for U.S. manufacturing overall. In the first column of Table 7, we report the average across all of manufacturing for the BLS TFP measure. With this as a benchmark, productivity in both textiles and apparel has surpassed the average for manufacturing. This is true for every census year in both industries for both the LP and BLS TFP measures.

4.3.3 Micro-Foundations of Productivity Growth: The case for creative destruction

When Schumpeter wrote of an economic structure “incessantly destroying the old one, incessantly creating a new one,” he could well have been thinking about the evolution of productivity dynamics. What does creative destruction imply about productivity dynamics? It implies that although many plants are simultaneously entering and exiting, the net impact of this churning is to increase productivity. Is this the case? Industry productivity, as typically defined, can increase either because surviving plants are actually becoming more productive or because market shares are being reallocated from less productive to more productive plants even when no single plant is becoming more productive. The former explanation seems more in line with an industry re-inventing itself. The latter explanation, on the other hand, is a bounded process in which plants aren’t actually changing, but rather the market is rewarding the more productive and penalizing the less productive. In order to assess the case for creative destruction, it is necessary to examine the micro-foundations of the industry-level productivity growth reported

in Table 7.

In order to investigate these questions and others, it is necessary to decompose the industry-level productivity figures and exploit the plant-level productivity estimates.²² Doing so requires some notation.

We define industry-level productivity, Ω as a weighted average of plant-level productivities. Hence, average annual industry productivity (reported already in Table 7) is given by:

$$\Omega_t = \sum_{i=1}^N s_{it} \hat{\omega}_{it} \quad (4)$$

where s_{it} is plant i 's weight in year t and $\hat{\omega}_{it}$ is the estimate of plant-level productivity. In keeping with the output-share weighted average used in Table 7, we use output shares as the weights in our decomposition. Our decomposition is given by:

$$\begin{aligned} \Delta\Omega &= \sum_{i \in C} s_{i,t-1} \Delta \hat{\omega}_{it} + \sum_{i \in C} \hat{\omega}_{i,t-1} \Delta s_{it} + \sum_{i \in C} \Delta s_{it} \Delta \hat{\omega}_{it} \\ &+ \sum_{i \in B} s_{it} \hat{\omega}_{it} - \sum_{i \in D} s_{i,t-1} \hat{\omega}_{i,t-1} \end{aligned} \quad (5)$$

Here, $\hat{\omega}$ is the deviation of each plant's productivity from the initial-year industry mean. The first term in the decomposition, $\sum_{i \in C} s_{i,t-1} \Delta \hat{\omega}_{it}$ represents the part of industry-level productivity growth due to within plant productivity growth. If the only reason industry productivity grew was that every single plant was becoming uniformly more productive, this term would capture all of the increase in productivity. Industry productivity might also increase, however, because output shares are being reallocated from less productive plants to more productive ones. This effect is often termed the reallocation effect, and it is captured by the second term in (12). The third term is harder to interpret. This term is a covariance term. It does not neatly fit into the rationalization versus real productivity taxonomy as it clearly contains both plant level productivity changes ($\Delta \hat{\omega}_{it}$) and reallocation (Δs_{it}). Our economic interpretation of this covariance term is that it is a measure of just how inter-related changes in plant-level productivity and changes in market share really are. To the extent that it is non-zero, one simply cannot separately consider productivity changes in the absence of reallocation. The fourth and fifth terms, respectively, capture the contribution of entry and exit. By construction, entry's contribution is positive and exits' is negative, but their sum is ambivalently signed. Entry and exit taken together might increase or decrease aggregate productivity.

²²There is an extensive literature on the best way to decompose industry-level productivity growth. See, for example, Foster, Haltiwanger, and Krizan (1998) and Levinsohn and Petrin (1999). We refer the interested reader to this literature for a discussion of the pros and cons of various decompositions.

We use this decomposition to examine what underlies the changes in aggregate productivity and, in particular, to examine evidence of creative destruction. Results are reported in Table 8. The results are pretty stark. The impact of reallocation is almost negligible. There is little evidence that, conditional on productivity levels, more productive plants are gaining output shares while less productive ones are losing shares. By far, most of the action is in within-plant productivity changes. When industry productivity increases, plants really are becoming more productive. This finding is striking in that it contradicts the oft-professed belief that in protected industries, there is little incentive for existing plants to become more productive. In textiles and apparel, it is within-plant productivity growth that is responsible for the increase in industry productivity.²³ The covariance term, which could be considered part of either reallocation or within-plant growth, is only moderate. The general message is one in which surviving textile and apparel plants are becoming more productive.²⁴

The last column gives the net effect of entry and exit. These figures address the question of whether all the plant entry and exit, on net, helps or harms industry productivity. In textiles, the net effect is always positive albeit modest. This is in part because entrants and exiters tend to have small market shares. In apparel, the magnitudes are also modest, although in two of the 5-year periods, the net impact of entry and exit was to (barely) decrease industry productivity.

The overall picture of productivity in the textile and apparel industries is one in which productivity increased over the two decades studied and in which most of this increase was due to plants truly becoming more productive. The net impact of entry and exit was consistently modest and usually contributed to higher industry-level productivity. We conclude that the decomposition of productivity seems to support the case for creative destruction. The fact that productivity increased mostly because of *within*-plant productivity gains strikes us as both surprising and hopeful. It is surprising because the tremendous job churning and plant entry and exit documented above might lead one to believe that one cannot teach an old dog new tricks—that it requires new jobs and new plants for productivity to increase. This result is hopeful because it suggests that incumbent plants can and do become more productive as they face increased competition both from abroad and from new domestic entrants. It also suggests that incumbent plants can learn to benefit from new technologies and new production techniques, although verifying this requires more analysis into the plant-level determinants of productivity and, in many cases, exit.

²³Of course, it could be that the within-plant productivity increases would have been yet more impressive in the absence of protection, but absent a well-specified model of the industries, this is speculation.

²⁴It is noteworthy that the only really large decline in within-plant productivity came in apparel during 1987-92. It might be that our productivity measure misses an important dimension of plant performance that is improving over this time period. After 1987, manufacturers were adopting quick-response production processes in order to accommodate the push toward lean retailing by retailers. The benefits of quick-response, which include improving the match between production and final demand and lowering inventories held by the manufacturer, might not be well reflected by our measure of productive efficiency.

4.3.4 Who Exits?

Before concluding our study of productivity, we want to examine more closely those plants that do *not* remain in our sample. The reason for this is that the characteristics of these plants speak to the issue of plain destruction versus creative destruction. If there is to be a benefit to the tremendous exit in these industries over the past two decades, it should be the case that exit is not a random phenomenon. This is an important part of the idea of creative destruction. On the other hand, if exit is both pervasive and random, it would be easier to make the case that these industries are simply declining. Section 2 makes clear that exit is pervasive. Here we examine whether it is random or whether instead exit seems consistent with loosely exposed notions of competition and industry evolution.

Table 9 presents the results of our analysis of what predicts plant exit in the textile and apparel industries. Many of the questions posed in the introduction under the title of “Interactions” are addressed here. The coefficients reported in Table 9 are the results of a probit regression. The dependent variable equals one if the plant exits and zero otherwise. The probits include fixed effects for 3-digit industry sub-codes. The coefficients in the table are the change in the probability of exit for an infinitesimal change in the independent variable for the case of continuous independent variables, and the change in the probability of exit for a discrete change in dummy independent variables. Each probability is evaluated at the mean value of the vector of independent variables. The coefficients on each regressor will be discussed in turn.

The first regressor is the Levinsohn-Petrin measure of productivity (hereafter productivity.) We find that, conditional on all the other regressors, a plant’s productivity explains exit in a very statistically significant manner. The negative coefficient indicates that as a plant becomes more productive, the probability of exit diminishes. This is the most important result in Table 9. Conditional on the size of the plant, the wages paid, the capital stock per worker, and measures of outsourcing, it is the less productive plants that are more likely to exit. There are multiple ways to interpret this result. Those whose faith in the forces of competition is such that they already knew this *had* to be the case can note that this then reflects well on our measure of productivity. Those who worry that the crazy-quilt of protection afforded by the MFA allows inefficient plants to prosper while protecting them from the realities of the world marketplace should find some solace in this result. In a probabilistic sense, inefficient plants die. One must be careful when interpreting the magnitude of these coefficients since the units in which productivity is measured, by themselves, have little meaning. Also, the units are not comparable across the two industries. Because the estimated coefficients are local derivatives (evaluated at the vector of means), we can derive a semi-elasticity by simply multiplying coefficients by the mean level of productivity in each industry. When we do this, we find that the semi-elasticity of the probability of exit with respect to a percentage change in productivity is -0.17 in textiles and -0.05 in apparel. This means that a 10 percent increase in productivity in textiles lowers the probability of exit by 1.7 percent conditional on all the other regressors and evaluated at

the mean. This effect is over three times as large as the analogous semi-elasticity for apparel.

The second regressor is the percent of work contracted out, our plant-level measure of outsourcing. In apparel, where outsourcing has been a defining change in the industry over our sample, we find that plants that contract out more work are less likely to exit, conditional on productivity and other plant-level characteristics. It has been frequently claimed that to succeed in the apparel industry, one needs to play the outsourcing strategy. This appears to be correct (at least in a probabilistic sense.) Evaluating all the regressors at their means, the magnitude of the coefficient for apparel implies that a ten percent increase in work contracted out lowers the probability of exit by about 1.6 percent. Given the very large increase over time in the percent of work that some plants contract out (much larger than the ten percent in the above example), this effect might be quite large.²⁵ In textiles, where outsourcing is much less of an issue, plants that outsource more are more likely to exit, conditional on all the other regressors. The production process in textiles is much less conducive to outsourcing, and it is not at all obvious that outsourcing is likely to be a successful business strategy. Those plants that do outsource may be doing so as a last resort. The magnitude of the coefficient implies that at means, a 10 percent increase in outsourcing in textiles increases the probability of exit by about 2.3 percent.

The third and fourth regressors are non-overlapping measures of in-sourcing. These variables, receipts for contract work and a dummy variable for whether the plant's main product class is commission work, capture the other side of outsourcing. For both textile and apparel, plants that do more insourcing have a higher conditional probability of exit. The positive sign on the coefficients on these regressors probably reflects the notion that these firms are taking on work that might otherwise be outsourced to foreign plants. In apparel especially, this is often the more labor-intensive part of the production process and it is not the segment of production in which U.S. plants are likely to have a comparative advantage relative to their foreign competitors. In apparel, plants that insource are competing with low-cost labor principally in Mexico and the Caribbean Basin. The results suggest that the U.S. plants that try to compete on this dimension are more likely to fail. The same phenomenon seems to be happening in textiles too, although both the magnitude of the coefficients and their statistical significance are smaller than in apparel.

"Percent skilled Labor" is a variable that gives the percent of the workforce classified as non-production workers.²⁶ Conditional on productivity and the out- and in-sourcing variables, and wages, a higher percentage of skilled labor is correlated with a greater probability of exit. Evaluated at the vector of means, a ten percent increase in the percent of skilled (i.e. non-production) labor implies that the probability of exit increases about 2.3 percent in textiles and about 1.5 percent in apparel. There are three interpretations of this result. The first is

²⁵The effect is not linear so it is not correct to simply extrapolate for large changes.

²⁶This is clearly not an ideal measure of skilled labor, but it is the best one can do with the data.

that, conditional on productivity, having more non-production workers is not a good idea; non-production worker labor is relatively expensive. The second is that before firms exit altogether, they may be more likely to first release the production workers— especially if production workers are paid hourly and work without annual contracts while non-production workers are paid a salary and may be more likely to have annual contracts. This is especially likely to be true if exit really means that the plant is moving operations to another location. Third, if there are fixed costs associated with hiring non-production workers (i.e. search costs), and a small chance that a plant with a temporary business slump may recover, then the option value of keeping non-production workers on payroll may exceed the direct cost.

“Inventories” is the percent of finished-good and work-in-progress inventories on hand at the end of the year relative to annual shipments. With the move to lean manufacturing, which emphasizes quick turnarounds, low inventory costs, and short product cycles, much has been written about how this strategy is key to success in these industries. Our plant-level results substantiate this claim. Even conditioning on productivity, larger inventories increase the probability of exit. Because our measure of productivity is unlikely to capture all (or even many) of the benefits of lean manufacturing, inclusion of inventories is unlikely to be highly co-linear with our measure of productivity. This effect is larger in apparel, where lean manufacturing is important for a broader set of products than in textiles. A ten percent increase in the ratio of inventories to shipments raises the probability of exit by 2.6 percent. The analogous figure for textiles is 1.6 percent.

We also find that, conditional on productivity, having more capital per worker raises the probability of exit in a statistically significant but quantitatively small way. The logic behind this is the same as that for non-production workers. Capital already enters the production function so our measure of productivity is taking into account the contribution of capital to output. Conditioning on productivity, having more capital raises the likelihood of failure. A simple explanation is that capital is expensive and having more than is productive is problematic. Additionally, the option value story described above for non-production worker labor may apply here as well; if there are fixed costs associated with building capital, a firm facing a downturn may hold onto its capital on the small chance that things turn around, even if most such firms will soon exit.

The next three regressors are measures of the cost of labor. Conditioning on all the other regressors, higher production worker wages increase the probability of exit in textiles while in apparel the effect is statistically insignificant. For non-production workers, higher wages are conditionally correlated with a lower probability of exit. This is consistent with unmeasured worker attributes that contribute positively to the success of the firm. In any case, these effects are of comparable magnitudes in textile and apparel and they are small. Increasing salaries by a thousand dollars lowers the probability of exit by about a tenth of one percent.

The last regressor is a control for plant size. Conditional on all the other regressors, it is

the case that big plants are less likely to exit. We interpret this as evidence of sunk costs in both industries. The finding that, regardless of almost any set of plausible regressors, bigger plants exit less frequently is a robust finding in the literature. This effect is somewhat large. An increase in plant size of 100 workers lowers the exit probability in textiles by 5 percent and in apparel (where a 100 worker change is a bigger deal as plants tend to be smaller) by 11 percent. Because textile plants are about twice as large as those in apparel, the analogous elasticities of exit with respect to employment are very similar in the two industries.

In conclusion, exit patterns seem to accord pretty well with what we might expect if the role of exit is to “dump the chumps.” Low productivity plants, those that insource and hence compete more directly with low-wage foreign competition, those that carry larger inventories, and those that pay high production worker wages are all more likely to exit. In this fashion, exit does indeed seem to be weeding out the less competitive plants, and this contributes to the process of creative destruction as opposed to just plain destruction.

5 Conclusions

Without a doubt, the U.S. textile and apparel industries have faced difficult times over the past quarter century. What is less obvious from the industry-level data is the process by which these industries are re-inventing themselves as they adapt to new technologies (in the case of textiles) and new organizational structures (in the case of apparel.) As we’ve documented above, there is substantial entry into the industries, job creation rates are high, and productivity dynamics suggest surviving plants have emerged all the stronger while it has been the less productive plants that have exited. To the question posed in the first line of the paper, the answer is that these industries are indeed examples of creative destruction. Although the industry-level evidence is certainly consistent with labeling the textile and apparel industries as declining industries, the plant-level evidence highlights substantial creation.

A natural question is whether intensified import competition can explain all of our results. We are not sure. Models of international competition and the evolution of distributions of industry outcomes over heterogeneous plants with simultaneous entry and exit are only now being developed. Recent work in this vein includes Bernard, Eaton, Jensen, and Kortum (2000) and Melitz (2000). The main results in this paper— simultaneous entry and exit, job creation and destruction, and changing distribution of plant-level productivities— provide a template for the sort of empirical phenomenon that these models need to be able to capture.

Even in the absence of appropriate models, at least one point should be clear: Use industry-level data with caution. Probably the definitive study of the apparel and textile industries and the role of international trade is Cline (1989). In his study (using industry-level data), Cline concluded that:

The persistence of protection in textiles and apparel, *despite its relative ineffectiveness in promoting adjustment* (especially in apparel) and its high cost to consumers reflects the perceived importance and political clout of the two sectors. (p. 6, Cline (1989)) (Emphasis added.)

We disagree. While we are silent on the cost to consumers of protection in these industries, plant-level evidence indicates tremendous adjustment.

Appendix A

In this appendix, we briefly discuss how we estimate productivity when employing the Levinsohn-Petrin approach.

Estimation proceeds in two steps. In the first step, we estimate the production function coefficients on those factors of production that are assumed to be freely variable and which will respond quickly to productivity shocks. In the second stage, we obtain the estimates of the coefficient on capital— a state variable in the plant’s optimization problem. We begin with a Cobb-Douglas production function for value-added²⁷

$$y_t = \beta_0 + \beta_k k_t + \beta_s l_t^s + \beta_u l_t^u + \omega_t + \eta_t, \quad (6)$$

where y_t is the log of value added in year t , k_t is the log of the plant’s capital stock, l_t^s is the log of the skilled labor input (number of non-production workers) and l_t^u is the log of the unskilled labor input (number of production worker hours). The two disturbance terms, ω_t and η_t comprise the productivity term and any measurement error on output. The first component of the productivity term contains the “transmitted” part of productivity, that is, the part that is observed by plants before they choose variable input levels and that may be serially correlated within a plant over time. The second component contains any i.i.d. differences in productivity that are not observed by plant before they choose input levels, as well as measurement error on plant-level output. The first component, therefore, is likely to be correlated with variable inputs.

In order to get consistent estimates of the coefficients on variable inputs, therefore, we must control for the effect that unobserved productivity has on plant choices of input levels. We use the observation in Levinsohn and Petrin (1999) that, with increasing marginal costs and a perfectly elastic demand for their good, plants will respond to an increase in productivity by using strictly more of variable inputs. This strictly monotonic relationship between productivity and variable input use allows us to write the unobserved productivity term as some exact function of the capital stock (on which variable input levels will also depend) and employment of a variable input:

$$\omega_t = h_t(e_t, k_t) \quad (7)$$

where h_t is the inverted function describing optimal energy use. We use energy, e_t , since it is excluded from the value-added production function.

With (5), we can rewrite (4) as

$$y_t = \beta_s l_t^s + \beta_u l_t^u + \phi(e_t, k_t) + \eta_t, \quad (8)$$

²⁷A variety of necessary conditions for the existence of a value-added production function are described in Bruno (1995)

where $\phi(e_t, k_t) = \beta_0 + \beta_k k_t + h_t(e_t, k_t)$. Equation (6) is partially linear; it is linear in variable inputs, and non-linear in electricity and capital. The goal in this first stage is to obtain estimates on the coefficients of the inputs that enter (6) linearly (i.e. β_s and β_u .) Here we follow the approach for semi-parametric estimation of a partially linear equation described in Olley and Pakes (1996). We project value added on l^s and l^u and on a fourth- or fifth-order polynomial expansion in e and k .²⁸ This allows us to obtain consistent estimates of coefficients on the included variable inputs (skilled and unskilled labor), and completes the first stage of the estimation routine.

In order to obtain estimates of productivity, we still require an estimate of the coefficient on capital. We estimate the β_k in our second step. Notice that capital enters $\phi(\cdot)$ twice, i.e. $\phi_t(e_t, k_t) = \beta_0 + \beta_k k_t + h_t(e_t, k_t)$, and hence β_k (and ω_t) are not identified without further restrictions.

Identification of the capital coefficient obtains from assuming that capital is quasi-fixed, and is slow to adjust to the productivity shock. While capital might well adjust to the expected part of a productivity shock, the identifying assumption maintains that capital does not instantaneously adjust to the unexpected part. To operationalize this notion, we must impose some structure on the stochastic process of the transmitted productivity shock, ω . Following Olley and Pakes (1996), we assume that ω_t follows a first order Markov process. In particular, we can write

$$\omega_t = E(\omega_t | \omega_{t-1}) + \xi_t, \quad (9)$$

where ξ_t is the “news” in the transmitted shock. To control for the endogeneity of capital, we only need to proxy for the $E(\omega_t | \omega_{t-1})$. From the first stage regression, we have an estimate of $\phi_t = \omega_t + \beta_k k_t$, a nonparametric function of which serves as our proxy for each plant’s expected productivity.

The moment condition that identifies the coefficient on capital is then given by:

$$E(\xi_t + \eta_t | k_t) = E(\xi_t | k_t) + E(\eta_t | k_t) = 0. \quad (10)$$

This moment condition simply states that capital does not respond to the innovation in productivity. Estimation proceeds by employing a Generalized Method of Moments estimator to find the parameter estimate that most closely matches our sample to the population moment analog.

²⁸Andrews (1991) shows asymptotic results for estimation of a partially linear equation (equation (6)) using a series expansion. Alternatively, one could use kernel-based methods to estimate a partially linear equation. See Robinson (1988). Levinsohn and Petrin (2000) uses this method, estimating a locally weighted quadratic least squares approximation. Readers not familiar with local regression smoothing might find it helpful to think of this step as using weighted least squares to construct predictions of y_t given (k_t, e_t) using a second order approximation in (k_t, e_t) (a quadratic in (k_t, e_t)). For any particular point (k_t^*, e_t^*) for which an estimate of the expected value of y_t is necessary, the regression weights the observations closest to the point (k_t^*, e_t^*) most heavily. A consistent estimator of $E(y_t | k_t = k_t^*, e_t = e_t^*)$ is then the intercept from this local quadratic regression. We have experimented in other contexts with the locally weighted least squares and found it gave, to a remarkable degree, the same results as the polynomial approximations.

Once we have estimated the coefficients of the production function, we then simply compute our plant-level measure of productivity, $\phi_{it} - \beta_k k_{it}$.

Table 1: U.S. Textile and Apparel Imports

	Textiles		Apparel		Total
	Imports	Exports	Imports	Exports	807 Imports
1972	3,155	1,340	3,640	596	179
1973	3,033	1,764	3,973	712	251
1974	2,977	2,123	3,876	972	387
1975	2,547	1,974	4,355	958	405
1976	3,062	2,200	5,639	1,102	445
1977	3,101	2,042	6,044	1,211	468
1978	3,612	2,196	7,646	1,436	577
1979	3,366	3,024	7,446	1,808	620
1980	3,507	3,306	7,631	1,887	639
1981	3,852	2,891	8,318	1,760	n.a.
1982	3,561	2,176	8,755	1,262	n.a.
1983	4,161	1,883	10,156	1,057	n.a.
1984	5,753	1,807	13,959	1,019	n.a.
1985	6,278	1,753	15,279	965	n.a.
1986	7,007	1,911	17,507	1,136	n.a.
1987	7,264	2,134	20,857	1,415	1,425
1988	6,602	2,530	20,834	1,808	1,730
1989	7,504	2,938	22,927	2,115	1,958
1990	7,499	3,706	23,372	2,490	1,993
1991	8,106	4,186	23,298	3,139	2,520
1992	8,802	4,487	27,448	3,862	3,202
1993	9,646	4,784	29,380	4,398	3,784
1994	10,379	5,292	31,879	4,960	4,468
1995	7,442	5,354	29,463	4,830	5,787
1996	7,424	5,674	30,008	5,283	6,434
1997	8,457	6,354	34,261	5,936	8,139

Values in million 1987 dollars.

Table 2: Establishments Count Data

	Apparel		Textiles	
	All Plants	Deleted ARs	All Plants	Deleted ARs
1972	24,441	19,724	7,203	5,611
1977	26,505	19,354	7,202	4,829
1982	24,391	21,054	6,630	5,071
1987	23,168	13,961	6,065	3,850
1992	23,093	14,379(*)	5,886	3,735(*)

(*) indicates estimates.

Table 3: Rates of Plant Entry and Exit

	Gross Rate of Entry	Gross Rate of Exit
<i>Textiles</i>		
1972-77	26%	32%
1977-82	31%	32%
1982-87	20%	38%
1987-92	28%	31%
<i>Apparel</i>		
1972-77	42%	44%
1977-82	48%	43%
1982-87	25%	55%
1987-92	49%	46%

Table 4: **Aggregate and Average Plant Employment and Capital Stock**

	Industry Employment (‘000s)	Mean Production Workers	Mean Non-Production Workers	Mean Real Capital Stock Per Employee
<i>Apparel</i>				
1972	1,368	77	13	\$8,041
1977	1,334	73	13	\$5,931
1982	1,189	63	14	\$5,935
1987	1,081	68	13	\$6,883
1992	985	62	13	\$6,493
<i>Textiles</i>				
1972	953	165	23	\$26,562
1977	875	173	26	\$24,807
1982	717	137	24	\$30,980
1987	672	151	26	\$32,893
1992	616	139	23	\$37,222

Table 5: **Average 5-year Rates of Gross and Net Job Flows**

	Production Workers		Non-Production Workers	
	Textiles	Apparel	Textiles	Apparel
Gross Job Growth	19	30	28	42
<i>Continuing Plants</i>	9	10	17	21
<i>Entrants</i>	10	20	11	21
Gross Job Loss	30	42	36	49
<i>Continuing Plants</i>	13	12	16	16
<i>Exiters</i>	17	30	20	33
Sum of Gross Job Flows	50	73	65	92
Net Job Flows	-11	-11	-6	-7

Table 6: **Outsourcing, Insourcing and Input Growth Rates**

	Plants	Average Production Worker Growth Rate	Average Non-Production Worker Growth Rate	Average Capital Stock Growth Rate
<i>Textiles</i>				
Outsourcers	4,819	-0.25	-0.24	-0.27
Others	18,520	-0.14	-0.07	-0.10
Insourcers	3,263	-0.23	-0.16	-0.16
Others	25,492	-0.15	-0.10	-0.14
<i>Apparel</i>				
Outsourcers	20,066	-0.14	-0.18	-0.18
Others	82,515	-0.14	-0.10	-0.13
Insourcers	59,533	-0.25	-0.23	-0.27
Others	74,693	-0.06	-0.04	-0.05

Table 7: **Productivity Indices**

	Manuf.	Textiles			Apparel		
	BLS TFP	LP Productivity	Value Added Per Worker	BLS TFP	LP Productivity	Value Added Per Worker	BLS TFP
1972	1.00	1.00	1.00	1.00	1.00	1.00	1.00
1977	0.98	1.11	1.09	1.12	1.09	1.41	1.07
1982	0.99	1.10	1.21	1.34	1.09	1.55	1.15
1987	1.10	1.58	1.71	1.46	1.38	1.74	1.22
1992	1.11	1.68	1.96	1.60	1.26	2.22	1.17

Table 8: **Decomposition of Industry Productivity Growth**

	<i>Percent of Total Change Due To:</i>			
	Output Share Reallocation	Covariance	Within-Plant Growth	Net Entry
<i>Textiles</i>				
1972-77	-3	13	67	12
1977-82	-6	13	-4	11
1982-87	0	9	79	10
1987-92	-1	11	58	16
<i>Apparel</i>				
1972-77	-5	21	38	20
1977-82	5	14	1	-9
1982-87	-3	22	49	14
1987-92	-2	5	-43	-2

Shares are averages over 3-digit industries.

Table 9: **Predicting Plant Exit**

	Textiles	Apparel
LP Productivity	-6.4818 (-9.20)	-1.7089 (-5.98)
Work Contracted Out (pct)	0.2309 (3.56)	-0.1622 (-7.12)
Receipts for Contract Work (pct)	0.0752 (3.27)	0.1644 (15.48)
Commission Work Dummy	0.0791 (5.72)	0.0837 (13.74)
Percent Skilled Labor	0.2386 (6.49)	0.1519 (8.13)
Inventories	0.1607 (4.00)	0.2569 (9.44)
Capital per Worker	0.0002 (2.52)	0.0005 (5.14)
Production Worker Wage	0.0021 (2.11)	0.0003 (1.18)
Non-Production Worker Salary	-0.0009 (-3.84)	-0.0011 (-7.47)
Other Compensation per Worker	-0.0064 (-7.18)	-0.0023 (-5.11)
Total Employment	-0.0005 (-24.72)	-0.0011 (-44.70)
Number of Observations	15,406	46,643

“z” statistics in parentheses.

References

- ABERNATHY, F. H., J. T. DUNLOP, J. H. HAMMOND, AND D. WEIL (1999): *A Stitch in Time*. Oxford University Press, New York.
- ANDREWS, D. W. K. (1991): "Asymptotic Normality of Series Estimators for Various nonparametric and semiparametric models," *Econometrica*, pp. 307–345.
- BALDWIN, J. R. (1999): "A Portrait of Entrants and Exits," *Statistics Canada, Analytical Studies Working Paper No. 121*.
- BERNARD, A., J. EATON, J. B. JENSEN, AND S. KORTUM (2000): "Plants and Productivity in International Trade," *NBER Working Paper No. 7688*.
- BRUNO, M. (1995): *NBER Macroeconomics Annual*. MIT Press, Cambridge and London.
- CLINE, R. (1989): *The Future of World Trade in Textiles and Apparel*. Publisher, City.
- DAVIS, S., AND J. HALTIWANGER (1992): "Gross Job Creation, Gross Job Destruction, and Employment Reallocation," *Quarterly Journal of Economics*, 107(3), 819–63.
- DAVIS, S., J. HALTIWANGER, AND S. SCHUH (1996): *Job Creation and Destruction*. MIT Press, Cambridge, MA, 1998 edn.
- FEENSTRA, R. (1996): "U.S. Imports, 1972-1994: Data and Concordances," NBER Working Paper 5515.
- FOSTER, L., J. HALTIWANGER, AND C. KRIZAN (1998): "Aggregate Productivity Growth: Lessons from Microeconomic Evidence," *NBER Working Paper No. 6803*.
- LEVINSOHN, J., AND A. PETRIN (1999): "When Industries Become More Productive, Do Firms?," *NBER Working Paper No. 6893*.
- (2000): "Estimating Production Functions using Inputs to Control for Unobservables," *NBER Working Paper No. 7819*.
- MARSCHAK, J., AND W. ANDREWS (1944): "Random Simultaneous Equations and the Theory of Production," *Econometrica*, 12(3-4), 143–205.
- MELITZ, M. (2000): "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Technology," Mimeo, University of Michigan.
- OLLEY, S., AND A. PAKES (1996): "The Dynamics of Productivity in the Telecommunications Equipment Industry," *Econometrica*, 64(6), 1263–1298.

OTA (1999): *Focus: An Economic Profile of the Apparel Industry*. Office Textiles and Apparel, U.S. Department of Commerce, Washinton, D.C.

ROBINSON, P. M. (1988): "Root-N-Consisten Semiparametric Regression," *Econometrica*, 56(4), 931-54.

RODRIGUEZ-ARCHILA, L. (2000): "Industry Trade and Technology Review," (3335).

SCHUMPETER, J. (1942): *Capitalism, Socialism, and Democracy*. Harper and Row, third edn.