Business Review

Federal Reserve Bank of Philadelphia

March • April 1995 ISSN 0007-7011

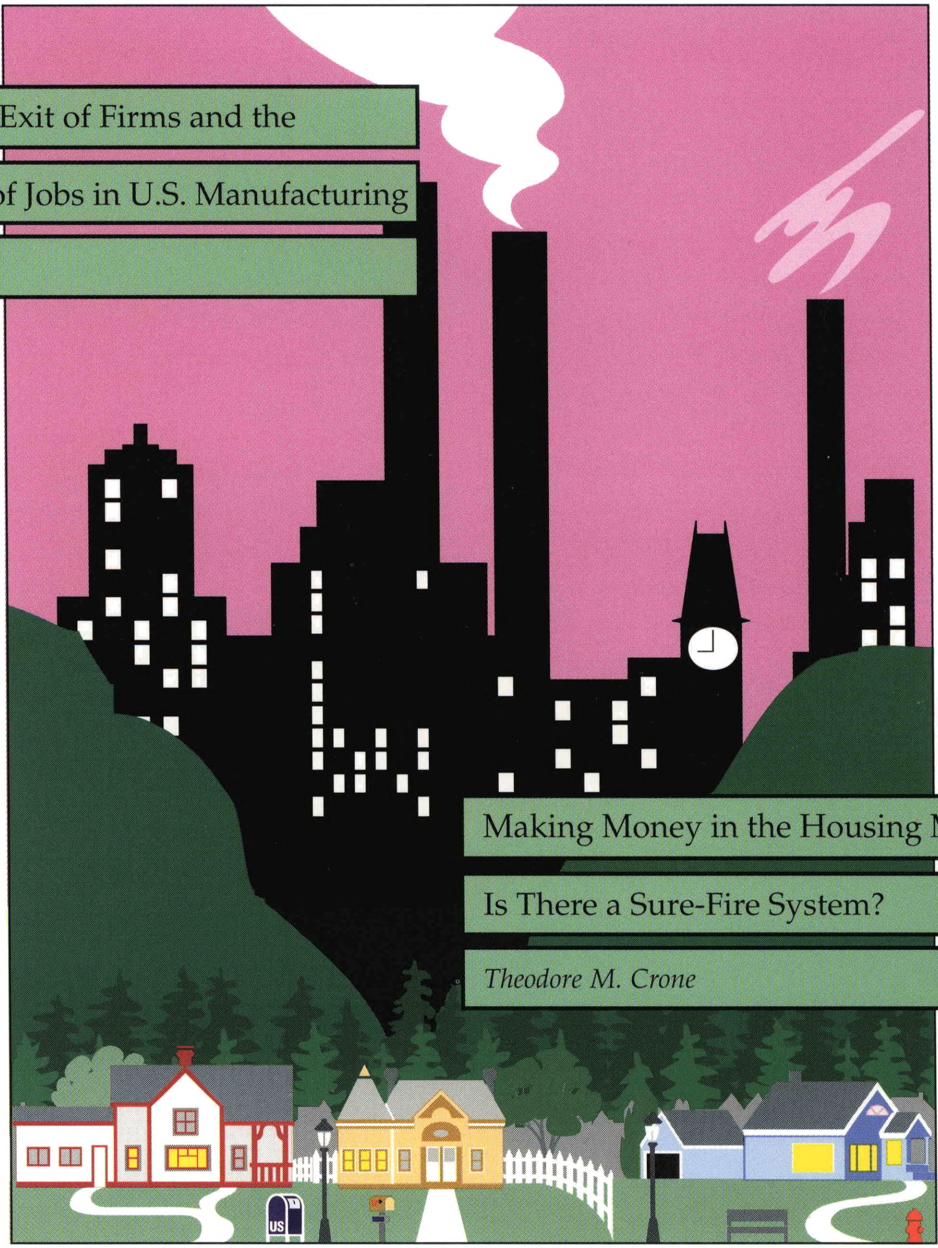

Entry and Exit of Firms and the Turnover of Jobs in U.S. Manufacturing

Rafael Rob

Making Money in the Housing Market: Is There a Sure-Fire System?

Theodore M. Crone

Business Review

The BUSINESS REVIEW is published by the Department of Research six times a year. It is edited by Sarah Burke. Artwork is designed and produced by Dianne Hallowell under the direction of Ronald B. Williams. The views expressed here are not necessarily those of this Reserve Bank or of the Federal Reserve System.

SUBSCRIPTIONS. Single-copy subscriptions for individuals are available without charge. Institutional subscribers may order up to 5 copies.

BACK ISSUES. Back issues are available free of charge, but quantities are limited: educators may order up to 50 copies by submitting requests on institutional letterhead; other orders are limited to 1 copy per request. Microform copies are available for purchase from University Microfilms, 300 N. Zeeb Road, Ann Arbor, MI 48106.

Please send subscription orders, back orders, changes of address, and requests to reprint to Publications, Federal Reserve Bank of Philadelphia, Department of Research and Statistics, Ten Independence Mall, Philadelphia, PA 19106-1574, or telephone (215) 574-6428. Please direct editorial communications to the same address, or telephone (215) 574-3805.

MARCH/APRIL 1995

Entry and Exit of Firms and the Turnover of Jobs in U.S. Manufacturing

*Rafael Rob**

The popular press and business and labor leaders have for quite some time been pronouncing the decline of U.S. manufacturing and the loss of jobs to overseas competitors. Aggregate data reveal that employment in U.S. manufacturing has, in fact, declined at an annual rate of 2 percent over the past 20 years. Behind this trend, however, lies a more intricate and interesting drama of industry dynamics. New firms constantly enter while others exit; some firms decline in size while others remain stable or grow; new jobs are created to replace some that are lost.

In this context, interesting questions arise as to the extent of the turnover process and the role of different types of manufacturing firms in this process. For example, what is the magnitude of entry into and exit from various industries? What are the characteristics of entering firms? Are they as large as incumbent firms? Are they as likely to survive? If they survive, do they grow as fast? In this article we survey the evidence that can be brought to bear on these issues and the theories that can rationalize the evidence.

A related issue that may have public policy implications is whether small firms are more efficient at creating employment opportunities than large firms. Supporters of this view consider small firms more dynamic and more

---

*Rafael Rob is a professor of economics, Department of Economics, University of Pennsylvania. When he wrote this article he was a visiting scholar in the Research Department of the Philadelphia Fed. The author acknowledges editorial comments from various members of the Research staff. He gives special thanks to Paul Calem for extensive suggestions. Any remaining errors are the sole responsibility of the author.

innovative, a claim which (to some extent) can be substantiated by the data: new firms, which tend to be smaller, are responsible for the creation of many new jobs, and many of these firms introduce new products that can potentially "take off" and generate yet more jobs. Opponents of this view, on the other hand, consider large and established firms as having proven themselves in the way they are managed or through the products and services they sell. Even though large firms are less likely to be innovative and to grow, they are also less likely to fail. Comparison of these views suggests that the relevant concept is the durability of jobs, not their mere creation. Another objective of this article, then, is to use the data reported here to construct a quantitative measure of the durability of jobs created by firms of different size.

THE EVIDENCE

The evolution of firms and industries is of interest to both business practitioners and economists. A business practitioner focuses on what accounts for the success or failure of individual firms and how firms can be profitably restructured, given the lessons learned from the experience of other firms. An economist focuses on systematic patterns that characterize the whole set of firms, for example, what the average lifetime of firms in the personal computer industry is and how it compares with that of the restaurant industry.

Although these questions arose in earlier literatures, interest in them has revived recently with the availability of more comprehensive data sets and improved methods of analyzing them, using fast digital computers.[1] The evidence in this section comes from two sources. The first is a series of papers by Timothy Dunne, Mark Roberts, and Larry Samuelson.[2] In these papers Dunne and his associates analyzed more than 300,000 manufacturing plants and the more than 200,000 firms operating them in the United States. These plants produce more than 99 percent of the output of 387 industries. Dunne, Roberts, and Samuelson followed these plants over a 20-year period and documented patterns of entry and exit, growth and decline, degree of diversification, and size distribution of firms, as well as how these variables are interrelated. The second source is a series of papers by Steve Davis and John Haltiwanger.[3] These papers address similar questions (with somewhat greater emphasis on macroeconomic issues), but they use a larger set of firms and more frequent observations.

Industry Turnover vs. Worker Turnover. Before proceeding it's important to stress both the connection and the distinction between

[1]Two early references are P.E. Hart and S.J. Prais in the United Kingdom and H. Simon and C. Bonini in the United States. Motivated by the striking similarity of the size distribution of firms (across time, industries, and countries), these researchers sought to explain the source of this similarity, to estimate the distribution in different industries, and to show the effect of public policies on it. Later contributions tried to relate the characteristics of various industries, such as the size of efficient plant, the advertising intensity, the degree of product differentiation, or industry growth, and their net rate of entry or exit. Their main goal was to understand how the nature of the industry or the product affects the turnover process and, in turn, industry structure and pricing behavior.

[2]Timothy Dunne, Mark Roberts, and Larry Samuelson, "Patterns of Firm Entry and Exit in U.S. Manufacturing Industries," *Rand Journal of Economics* 19 (1988), pp. 495-515; "Plant Turnover and Gross Employment Flows in the U.S. Manufacturing Sector," *Journal of Labor Economics* 7 (1989a), pp. 48-71; "Firm Entry and Postentry Performance in the U.S. Chemical Industries," *Journal of Law and Economics*, 32(2), Part 2, 1989(b), pp. S233-71.

[3]Steve Davis and John Haltiwanger, "Gross Job Creation and Destruction: Microeconomic Evidence and Macroeconomic Implications," *NBER Macroeconomics Annual*, 5 (1990), pp. 123-68; "Gross Job Creation, Gross Job Destruction and Employment Reallocation," *Quarterly Journal of Economics* 107 (1992), pp. 819-63.

industry turnover and worker turnover. The connection arises because when one plant shuts down and another plant opens, some jobs are lost and others are created, so this generates reallocation of the labor force. In this way, firm turnover gives rise to worker turnover.

But there are other reasons workers look for new jobs. By looking for a new job a worker can create a better match between his or her qualifications and the job's requirements. Furthermore, because it takes time to determine the quality of a match and because new workers constantly arrive over time, workers are continually in the process of sorting and resorting themselves into more suitable jobs.

Another reason behind worker turnover is that workers accumulate knowledge and experience at their present jobs, qualifying them for more advanced positions. Yet, such positions are not always available with their present employer, either because of the nature of their present employment or because others are already occupying the more advanced positions. Thus, a certain fraction of job switches occurs for career or advancement reasons.

Finally, some switches are induced by temporary changes in the conditions that firms face: workers are laid off from jobs during bad times, with the possibility of recall at a later date when business conditions improve. By that time, however, some of these workers have found different jobs, and consequently, their positions are filled by others. Given these causes, workers reshuffle across jobs, although the jobs themselves remain intact. Thus, worker turnover is not perfectly matched by firm turnover.[4] One estimate states that the fraction of job switches induced by firm turnover is between 35 and 56 percent.[5]

Volume of Turnover. Let's turn to the features of the data. The most striking finding is the magnitude of firm turnover: every year an average of 8 percent of all incumbent firms in manufacturing exit and an average of 9 percent enter, resulting in 17 percent turnover with net entry of 1 percent.[6] Likewise, the rate of job destruction is 11 percent a year, and the rate of job creation is 9 percent, resulting in 20 percent turnover with a net loss of 2 percent.[7] The entrants represent either de novo firms (a new firm with a new production facility) or diversification by a firm already operating in another industry but now changing the mix of outputs in its plant or adding a new plant. The breakdown between these two categories (averaged over all manufacturing industries) is 55 de novo firms to 45 non-de novo.[8]

Variation Across Industries. While these numbers represent averages across all firms in the manufacturing sector of the U.S. economy, there's a great deal of variation between industries. Some broad industry categories—for example, lumber and apparel—exhibit high firm and job turnover; others—for example, tobacco and primary metals—exhibit low turnover.[9] Usually, industries that show high entry rates (rates refer to the *gross* rate unless other-

[4]To distinguish between the two I will refer to changes that emanate from firms as job creation and destruction (or job turnover) and to changes that emanate from workers as employment search; the fact that a worker changed her job (for whatever reason) is referred to as a job switch.

[5]Davis and Haltiwanger, 1992, pp. 833-84.

[6]See Dunne, Roberts, and Samuelson, 1988, p. 503.

[7]Davis and Haltiwanger, 1992, p. 820.

[8]See Dunne, Roberts, and Samuelson, 1988, p.504.

[9]Industries are often classified as two-digit and four-digit, referring to the Standard Industrial Classification (SIC) of sectors in the U.S. economy. The two-digit classification is a coarse breakdown of sectors into 20 categories numbered from 20 through 39, the tobacco sector corresponding to 21, for example. The four-digit classification is a refinement into subsectors, the cigarette subsector corresponding to 2117, for example. Further details on this classification system can be found in F.M. Scherer, *Industrial Market Structure and Economic Performance* (Houghton Mifflin Company, 1980).

wise specified) simultaneously show high exit rates, implying high turnover. Industries that exhibit a higher than average entry rate in a given year tend to exhibit a higher than average entry rate in the following year. The extent of turnover is a characteristic of the industry that tends to persist over time.

The degree to which industrial turnover varies across industries can be measured in gross rates of entry and job creation. Firm entry rates range between 5 and 13 percent across industry categories, the average being 9 percent. Gross job creation rates range between 5.9 percent in the tobacco industry and 12.9 percent in the lumber industry (Table 1).[10] Furthermore, even within a given industrial category, there's a great deal of variation across subcategories. For instance, while the average entry rate for the food-processing industry is 8.9 percent, 10 percent of its subindustries exhibit an entry rate not exceeding 1.4 percent, while 10 percent exhibit an entry rate above 15 percent.

Net vs. Gross Entry and Job Creation. Since industries with high entry rates also usually have high exit rates, net entry of firms bears little relationship to gross entry and turnover. Take, for example, the transportation industry, in which the number of firms declined only 0.2 percent per year over the sample period. Yet, it experienced gross rates of entry and exit of 9.1 and 9.3 percent, respectively, each of which is almost 50 times larger than the net change![11] Similarly, the gross job creation rate in the transportation industry averaged 9.4 percent per year, and job destruction averaged 9.9 percent, both much larger than net job creation (Table 1). Thus, we cannot view turnover in manufacturing as if it were stemming from sectoral shifts alone, that is, firms and jobs moving from declining industries into growing industries. On the contrary, the data show that even within the same industry, large flows of entry and exit occur, which shows that turnover is based as much on the characteristics of the firm as on the environment in which it operates.

Another piece of evidence supports this fact: the magnitude of turnover persists across all phases of the business cycle. Even 1975, the worst downturn year between 1973 and 1986, saw an entry rate of 6.7 percent, while 1973, the best upturn year, saw an exit rate of 6.1 percent.[12,13] Similarly, even in average years there is substantial entry and exit. Hence, it's unlikely that macroeconomic changes alone drive the turnover process, that is, firms exit when conditions are bad and enter when they improve. Instead, the characteristics of individual firms (or changes in these characteristics) are important in explaining how macroeconomic changes will affect firms.

Persistence and Concentration. The turnover process demonstrates both persistence and concentration. When a firm enters, it's more likely to stay subsequently than to exit, and the likelihood of continued survival rises over time.[14] Similarly, when a diversified firm exits an industry, it's more likely to remain inactive thereafter than to resume activity. Also, when an incumbent firm grows, it's more likely to retain its size than to decline again. Therefore, these events reflect a *persistent* change in a characteristic of a firm.

Second, changes in employment affect certain firms much more than others. For instance, firms that fire more than 50 percent of their workers account for 34 percent of job

[10]Dunne, Roberts, and Samuelson, 1988, p.505.

[11]Dunne, Roberts, and Samuelson, 1988, p. 506.

[12]Davis and Haltiwanger, 1992, p. 830.

[13]One curious feature of the data is that entry is countercyclical, i.e., it tends to be relatively high in downturn years.

[14]Dunne, Roberts, and Samuelson, 1988, p. 510.

TABLE 1

Gross and Net Job Creation Rates by Industry

(size-weighted annual average[a] percent change)

Industry (SIC code)	Job Creation	Job Destruction	Net Job Creation	Job Turnover	MAX
Food (20)	8.9	10.4	-15.0	19.3	10.8
Tobacco (21)	5.8	8.2	- 2.4	14.0	9.0
Textile (22)	7.4	11.0	- 3.6	18.5	12.4
Apparel (23)	11.6	15.6	- 4.0	27.2	16.8
Lumber (24)	12.9	16.0	- 3.1	28.8	18.8
Furniture (25)	10.1	12.1	- 1.9	22.2	14.3
Paper (26)	6.3	7.8	- 1.5	14.1	8.9
Printing (27)	9.1	8.7	+ 0.4	17.8	9.9
Chemicals (28)	6.8	8.0	- 1.3	14.8	8.9
Petroleum (29)	6.6	9.1	- 2.5	15.7	10.0
Rubber (30)	10.7	11.8	- 1.1	22.5	14.3
Leather (31)	9.1	14.4	- 5.3	23.5	15.2
Stone, clay and glass (32)	9.3	12.3	- 3.1	21.6	13.6
Primary metals (33)	5.9	11.4	- 5.4	17.3	12.6
Fabricated metals (34)	9.5	12.0	- 2.5	21.5	13.7
Nonelectric machinery (35)	9.6	12.1	- 2.5	21.7	14.1
Electric machinery (36)	9.7	10.9	- 1.1	20.6	13.0
Transportation (37)	9.4	9.9	- 0.6	19.3	12.3
Instruments (38)	9.3	9.3	- 0.2	18.6	11.2
Miscellaneous (39)	10.8	14.5	- 3.7	25.3	15.6
Total manufacturing	9.2	11.3	- 2.1	20.5	12.9
Size-weighted cross-industry standard deviation	1.6	21.0	1.5	3.4	2.3

[a] Size-weighted average based on annual values for 1973-86 (but not 1974, 1979, 1984)

Table based on data from Table II in Steven Davis and John Haltiwanger, "Gross Job Creation, Gross Job Destruction, and Employment Reallocation," *Quarterly Journal of Economics*, 107 (1992), p. 831, excerpted with permission from The MIT Press, Cambridge, MA.

destruction.[15] Therefore, job destruction is concentrated in some firms instead of being spread evenly across firms, showing again the effect of firm-specific attributes.

These findings confirm that heterogeneity across firms is a crucial feature of the turnover process. Individual firms may still be affected by outside factors—for example, sectoral shifts and business fluctuations—but whether a particular firm can weather these fluctuations or even prosper from them depends on its innate characteristics: how well it is managed, what labor relations within it are like, whether it is

[15]Davis and Haltiwanger, 1992, p. 836.

innovative, and so forth.

The Size of Entering and Exiting Firms. The next issue is, what observable attributes characterize firms that enter or exit? The most apparent attribute is size: although 8 to 9 percent of firms turn over in any given year, entering firms account for only 3 percent of the output in the industry they belong to. Therefore, the size of entering firms is significantly smaller than the size of an average incumbent firm. Dunne and associates estimate the average output of an entrant to be 35.2 percent of the average output of an incumbent firm.[16] A similar pattern is detected for exiting firms: they are smaller than the average surviving firm.[17]

Patterns That Arise Over Time. Next, let's consider how newly established firms change over time, particularly their chances of surviving and growing. On average, across industry categories, the market share of entrants during their first five years is 16.2 percent; in their next five years it's 10.4 percent; then 8.3 percent.[18] Therefore, if we follow a set of firms that enter in the same year (also known as a cohort), we see that their market share steadily declines over the years. This represents two opposing forces. First, during the three consecutive five-year periods, the size of the average surviving firm rises from 35.2 percent to 54.3 percent to 127 percent of the average firm size in its industry. On the other hand, the share of firms surviving decreases with age from 38.5 percent to 19.9 percent to 14.0 percent. Therefore, while young firms increase in size if they survive, their chances of survival are smaller. On balance, the latter effect dominates the former, resulting in a decreasing market share as a cohort ages. A second prominent feature is that the standard deviation of survival rates and market shares declines as a cohort ages. That is, over time the members of a cohort become more homogeneous, and uncertainty about their future prospects is diminished.

Hence, small firms partake more actively in the entry and growth process, but they are less likely to succeed and stay around for a long time.

JOB CREATION: SMALL VS. LARGE PLANTS

The turnover processes we have been examining—entry and exit of firms and survival, growth, or demise of entrants—are important components of job turnover overall, but they are not the only components. Further contributing to the creation and destruction of manufacturing jobs are growth and decline among well-established firms. This section combines information on entry and exit with information on growth or decline of existing firms to determine the net impact of all of these dynamic influences. In particular, we compare firms in different size categories and look at their contributions to job stability, taking into account the durability of jobs as well as their creation.

To that end, consider Table 2, which is based on annual observations of manufacturing plants during the period 1973-88.[19] The

[16]Dunne, Roberts, and Samuelson, 1988, p. 503.

[17]These numbers pertain to the entire population of entrants. De novo entrants account for 55.4 percent of the number of new firms, but for only 50 percent of the output of new firms. On the other hand, diversifying firms with new plants account for 8.5 percent of the number and 14.4 percent of the output of all new entrants. The size of a de novo entrant is only 28.4 percent of the average firm size in its industry, while the corresponding size of a diversifying firm is 87.1 percent. Thus, diversifying entrants tend to be larger than de novo entrants.

[18]Dunne, Roberts, and Samuelson, 1988, p. 510.

[19]Data in Table 2 are taken from *Quarterly Journal of Economics*, 107, 1992, Davis and Haltiwanger, "Gross Job Creation, Gross Job Destruction, and Employment Reallocation, p. 841, and excerpted with permission of The MIT Press, Cambridge, Massachusetts.

TABLE 2
Gross Rates of Job Growth and Decline by Plant Size
(annual percent change)

No. of employees	1-99	100-249	250-499	500-999	1000+
POSITIVE	14.0	9.9	8.6	7.0	6.0
NEGATIVE	16.4	12.0	10.5	9.3	7.8
SHARE (%)	24.6	18.5	16.2	13.4	27.3

columns of this table correspond to the size of a plant at the date of observation; for example, 100-249 corresponds to a plant employing between 100 and 249 workers. The row labeled POSITIVE shows, for each size category, the average rate of job growth across all observations of plants that happen to be growing (which includes new entrants or growing establishments). NEGATIVE shows, for each size category, the average rate of job loss across all observations of plants that are declining in size. For example, plants that employ between 100 and 249 people and were declining in size experienced an annual rate of job loss of 12 percent, on average. The last row in this table (SHARE) represents the share of the industry employment accounted for by different-sized plants. For example, plants with more than 1000 employees accounted for 27.3 percent of employment in the manufacturing sector. This table shows that the rates of growth and decline are both smaller for larger plants, which shows the greater stability of these plants, that is, large plants create jobs less rapidly, but they also lose jobs less rapidly.

These numbers represent *averages* over all plants. *Individual* plants, however, may experience different rates of growth or decline—even if they belong to the same size category. For example, while the average growth rate in the positive category for the 100-249 size group was 9.9 percent, some plants may have increased by only 5 percent, while others increased by 15 percent. Likewise, if we take a long-term perspective, different plants may undergo dramatically different employment histories—even when they start out with the same number of employees. To give a simple numerical example, a certain plant may have employed 45 workers in its first year in business, then 49 in the second year, then 32, 47, and 15 in the third, fourth, and fifth years, and then may have gone out of business by the end of its fifth year. Altogether, over the time it was operating, such a plant provided 188 annual jobs. Another plant may have undergone a different employment history: 45, 22, and 0 (going out of business by the end of the second year), resulting in a total employment of 67 annual jobs. Therefore, although these plants started out with the same number of employees, they stayed in business a different number of years and provided a different number of jobs every year they operated. For ease of reference we shall call the total number of annual jobs generated over the time a plant was operating "job years."

In the example above, we simply added up jobs in different years without considering the timing of job creation. But a firm that starts with 100 employees and contracts to 80 employees in its second year can be viewed as generating jobs of greater total value during those two years than a firm that starts with 80 people and a year later expands to 100, because the first firm is generating jobs sooner. In this view, which we adopt, the sooner jobs are created the better, holding total annual jobs constant. Thus, employing a discount factor of 4 percent, we shall discount jobs in later years in counting the jobs generated over a particular firm's employment history.[20] If we apply this rate to the two examples in the previous paragraph, we obtain 176 and 66 discounted job years (instead of 188 and 67).[21]

Contributions of Firms in Various Categories to Job Stability. The next task is to determine the number of jobs that the typical firm in a given size category can be expected to generate over future years. This will yield comparisons between firms in terms of how effectively they create durable job opportunities. This task is accomplished by using information about the growth and decline of firms' employment by size category (as shown in Table 2) to generate a numerical assessment of the various possible employment histories and the resulting average number of discounted job years. The appendix spells out the technical details of this estimation procedure.

The first row in Table 3 (national average) shows the results of this estimation (ignore, for now, the remaining rows). As the table shows, the average number of future jobs generated by a representative firm steadily increases with its current size, that is, a firm that is large today is expected to provide more jobs over its lifetime than a firm that is small today. This finding should come as no surprise. First, a large firm is providing more jobs at present. Second, large firms' persistence rate is higher, that is, a large firm is more likely to retain its size than to decline (and a small firm is more likely to exit than to stay in business). Both factors favor large firms as generators of durable employment opportunities. The problem with the argument in favor of small firms is that it puts too much weight on the increase in their size when they happen to succeed, overlooking the large number of small firms that fail. As the data show, a small firm is more likely to fail (on average) than a large one, a fact mirrored by the smaller number of discounted

[20]I chose this rate because it represents the rate by which investors might discount riskless future earnings. The results reported below have the same qualitative feature for interest rates between 0 and 20 percent.

[21]The discounted total jobs in the first example equal 45 + 49/(1.04) + 32/(1.04)^2 + 47/(1.04)^3 + 15/(1.04)^4 = 176.31. Those in the second example equal 45 + 22/(1.04) = 66.15.

TABLE 3

Discounted Job Years Generated by Different Sized Firms Nationwide and in Different Regions

No. of employees	1-99	100-249	250-499	500-999	1000+
REGION					
National Average	2402	3661	5881	9541	15,626
New England	2310	3552	5783	9501	15,817
Middle Atlantic	1969	2947	4772	7962	13,784
South Atlantic	1982	3023	4974	8397	14,657
E. South Central	2390	3681	5979	9768	16,113
W. South Central	2536	3954	6443	10,483	17,098
E. North Central	2360	3636	5915	9688	16,041
W. North Central	2591	3921	6220	9899	15,829
Mountain	3060	4608	7161	11,042	16,920
Pacific	2662	3929	6089	9497	14,904

job years a small firm can be expected to generate. A recent paper by Steve Davis, John Haltiwanger, and Scott Schuh confirms the result that small firms generate less durable jobs.[22]

Put another way, the results in Table 3 show that while small firms tend to grow, on average, relative to their initial size, enough of them fail so that the representative small firm will not create as many durable jobs as the representative large firm that exists today.

A similar procedure can be used to assess the number of discounted job years generated by firms of different sizes in different regions of the country. The results of this estimation are shown in the last nine rows of Table 3. The pattern revealed in these rows is similar to the national pattern: a large firm can be expected to generate more lifetime jobs than a small one. However, the differences between large and small firms are starker for some regions. For example, in the South Atlantic and the West South Central regions, the "handicap" of small firms compared with large ones was bigger than that in the Pacific and Mountain regions. Therefore, the long-term prospects of small firms were better in the western regions.

Next we consider the effect of firms' age: is a young firm more likely to generate more discounted job years than an older firm, or the other way around? Using a procedure similar to the one used to generate Table 3, we estimated the number of discounted job years likely to be created` by firms of different ages (Table 4). This table reveals an interesting pattern: the age effect for small plants is negative, that is, the older a small firm is, the fewer discounted job years it is likely to generate. For successively larger firms, the age effect is still negative, but it's weaker. Finally, for the largest firms the age effect is positive. Therefore, we have a positive interaction between size and age: the number of discounted job years increases with firm size, but it increases even faster if we allow a simultaneous increase in firm age. One possible interpretation is that if a firm is small and old, chances are it is "on its

[22]See Steve Davis, John Haltiwanger, and Scott Schuh, "Small Business and Job Creation: Dissecting the Myth and Reassessing Facts," National Bureau of Economic Research Working Paper 4492, October 1993.

TABLE 4

Discounted Job Years Generated by Firms in Different Age Groups

No. of employees	1-99	100-249	250-499	500-999	1000+
AGE IN YEARS					
1	3473	4764	6776	9674	13,813
2	2363	3439	5331	8418	13,544
3	2741	4305	7000	11,277	18,079
4-5	2261	3457	5617	9236	15,437
6-10	2235	3496	5798	9696	16,452
11-14	2058	3278	5574	9598	16,908
15+	1642	2642	4673	8599	16,958

way out." On the other hand, if a firm is large and old, its method of operation or its product enables it to generate large sales and to survive for a long time. Thus, age provides supplemental evidence concerning the success of firms, although the implications are asymmetric across small and large firms.

Finally, we can analyze the effect of firm structure, particularly the employment prospects of single- vs. multiple-plant firms. The results of the same estimation procedure for this case are shown in Table 5, which demonstrates that a multiple-plant firm can be expected to generate a larger number of jobs than a single-plant one. However, the advantage of being part of a multiple-plant firm declines as a plant increases in size. The ratio of discounted job years generated by multiple-plant vs. single-plant firms declines as plant size rises. Thus, the feedback between size and multi-plant status is negative. A possible interpretation is that a plant that's part of a multiplant firm can freely "borrow" the expertise of its parent company, giving it an advantage over a firm that has no access to such expertise. However, once a plant reaches a large enough size, it is successful enough on its own, and the ownership effect is less relevant.

While these results provide support in favor of large firms as generators of employment, one should be careful in using them for policy analysis. The numbers in the tables reflect only the discounted number of job years generated by representative firms of different sizes, not the differential costs of operating these firms or the subsidies that might be needed to sustain them at their present size or to cause them to grow (which is much harder data to come by). It's quite conceivable that the subsidy needed to create a new job at a large firm is higher than the corresponding subsidy for a small one. Whether a given subsidy or tax break can stimulate more new jobs at a large firm will depend on the firm's effectiveness in creating an extra job, which is a separate issue from the durability of a job once it is created. Hence, further empirical analysis is needed to determine the effectiveness of subsidies in the hands of large vs. small firms (and to balance that against the differential durability of jobs).

TABLE 5
Discounted Job Years Generated by Single- vs. Multiple-Plant Firms

No. of employees	1-99	100-249	250-499	500-999	1000+
OWNERSHIP TYPE					
Single-plant firm	1927	2966	4935	8435	14,950
Multiple-plant firm	2978	4449	6888	10,619	16,314
Difference (multi - single)	1051	1483	1953	2,184	1,364
Ratio (multi/single)	1.545	1.50	1.395	1.258	1.091

THEORIES THAT EXPLAIN FIRM TURNOVER

The sections above surveyed the facts about industry dynamics and provided a method for estimating the number of jobs generated by representative firms of different sizes, but they did not elaborate on the basic forces that drive the turnover process. The purpose of this section is to survey such theories and indicate how they relate to the systematic patterns shown by the data.

Broadly speaking, theories of firm turnover fall into three categories: passive learning, active learning, and adjustments to outside disturbances.

Passive Learning. This theory is based on the premise that firms find out about their suitability to an industry (that is, their relative efficiency) only by operating in it.[23] According to this theory, whether a firm belongs in an industry is an inherent characteristic (a type) that remains unchanged over time but that can be discovered only through experience. Therefore, the process of entry and exit can be thought of as natural selection or survival of the fittest. This idea may explain the fact that young firms have a comparatively low survival rate but also a comparatively high growth rate when they survive. This follows from the fact that young firms are largely uncertain about their type. Once they have operated, they learn about their type and the uncertainty is reduced. As a consequence, either such firms become dismayed and leave or they receive favorable information and are able to grow rapidly. This also explains why the variability of growth rates is highest among young firms and why it declines with age, as shown by the evidence cited earlier.

This model explains other historical patterns not previously discussed, including the positive correlation between size and profits. That is, efficient firms produce large quantities and generate large profits at the same time. Likewise, as the industry matures, it becomes more concentrated as more efficient firms gain market share at the expense of less efficient firms. Accordingly, this process produces a positive correlation between concentration and average profitability. Also, this process induces a positive correlation between concentration and variability of profits. Again, this occurs because a firm's size and profits are directly related to the firm's efficiency, and efficiencies diverge over time as firms are sorted out.

The main shortcoming of this theory is that firms in the real world continually enter and exit even in mature industries, but this theory predicts such a process should eventually subside (unless industry demand keeps growing). Nonetheless, this theory has retained its popularity because it can explain a broad range of systematic patterns. It seems especially relevant to industries in which the success of firms depends on a difficult-to-alter specialized asset (a manager or a particular location or raw material).

Active Learning. The basic premise that distinguishes this theory from passive learning is that a firm's type, that is, its suitability to a given industry, changes during its tenure in the industry. This change may be the result of any number of things: successfully completing a research and development project, developing a new product and successfully marketing it, hiring a particularly successful manager, or raising morale among its employees. Alternatively, all or some of these endeavors may fail, leading to an unfavorable change in the firm's type. In some formulations of the theory the process by which a firm's type changes is explicitly incorporated, while in others only the net outcome of such a process is specified.[24,25] Either way, an important consequence of the active learning premise is that firms

[23]See Jovanovic, 1982, for a comprehensive treatment of the theory.

continually enter and exit even though the industry remains stable over time.

The main findings of this theory are as follows. First, if the process of learning to be successful exhibits persistence—that is, if currently efficient firms are expected (on average) to remain efficient in the future—then, for a given cohort of firms, the range of firms' sizes, profits and stock-market valuations, and the likelihood of firm survival all increase over time.

Second, increased cost of entry lowers the rate of both entry and exit, the rate of turnover, and the number of operating firms. Therefore, the industry becomes more sluggish, and firms collect returns on their cost of entry over a longer period. This lowers the average profitability in the industry but raises the profitability and market share of larger and more efficient firms. It also produces a positive correlation between the two, which we historically observe.[26]

Third, increasing the demand for an industry's output raises the entry rate as well as the number of firms in the industry but leaves all life-cycle properties—such as longevity and probability of survival—intact. Furthermore, the effect that increasing demand has on output, prices, and wages depends on how inputs are priced: if inputs command the same price no matter how many are purchased, the output increases but without an increase in the product price; otherwise, both the product price and real wages increase, the degree to which they increase depending on conditions in the labor market.

Fourth, a higher fixed cost makes it more costly for a firm to stick around during "bad times" (for example, when the variable cost of production is high), which leads to more firm attrition and an increase in the efficiency of operating firms and in their profits. This feature is the theoretical counterpart of how the presence of economies of scale leads to larger firms.[27] The main shortcoming of this theory is that it does not adequately explain what causes changes in firms' efficiencies and how different sources of change affect firms differently.

Adjustments to Outside Disturbances. The premise behind this theory is that firms enter or exit, or grow or decline, in response to outside changes (called shocks) that affect their industry.[28] For instance, firms respond to changes in the demand for their final product (because of changes in consumer tastes, for instance), the costs of entry and exit, or the prices they have to pay for inputs. Consequently, the entry and exit process can be understood as mirroring these changes and how firms react to them. One feature of this theory is the influence of history: a firm that enters in response to a favorable demand disturbance will not later exit if the disturbance merely reverses itself; to induce exit, a larger negative disturbance is needed. Therefore, how many firms operate in an industry at a given point depends on the history of demand disturbances, not merely on the present state of the industry.

This theory also establishes that a higher cost of entry will reduce the rate of turnover and increase the variability of the value of firms. This is consistent with what is observed in sectors such as agriculture or housing (where there is a large sunk investment): the value of firms fluctuates over a much wider range than in retail trade or restaurants (where there is a small sunk investment). This occurs because

[24]For example, Ericson and Pakes, 1990.

[25]Hopenhayn, 1992.

[26]Similar features arise if the cost of entry is the same, but the productive efficiency is worse upon initial entry into the industry.

[27]See Orr, 1970.

[28]See Lambson, 1992.

the turnover process in industries with high sunk cost is sluggish, and it's more economical for firms to weather changes than to exit and re-enter.

Furthermore, the theory attributes the coexistence of entry and exit to infrequent changes in market conditions. For instance, the oil shock of the 1970s induced exit of energy-inefficient firms, development of more energy-efficient techniques, and entry of new firms that use these techniques. The limitation of this theory is that it relies on such infrequent shocks to explain simultaneous entry and exit, while the data show that entry and exit occur on a regular basis without the occurrence of any unusual disturbances.

CONCLUSION

This article argues that the entry and exit of firms is not solely driven by factors external to them—macroeconomic changes or sectoral shifts that drive firms out of declining industries and into growth industries. Instead, the data show that entry of new firms occurs even in the worst downturn years and exit occurs in the best upturn years. Similarly, even when the number of firms in an industry declines, new firms still enter and grow. Therefore, to account for entry and exit one should look at characteristics of *individual* firms—size, age, whether they operate in a single industry or several, where they are located, and so on.

Considering these characteristics allows us to evaluate the number of jobs different types of firms create. New firms tend to be smaller, and they create the bulk of new jobs. On the other hand, small firms tend to be short-lived, and hence, so are the jobs they create. Therefore, to compare small and large firms, this article introduced a unit of measurement that combines the flow of new jobs and their persistence. We compared small and large firms in terms of the number of durable job opportunities they generate.

Using a numerical exercise, we showed that a representative large firm can be expected to generate a larger number of durable job opportunities because of the greater stability of these firms. In other words, a large firm has a smaller probability of decline, so once it has created a job, it tends to persist. This finding holds for firms in different regions of the country, although the effect is somewhat weaker in the higher turnover regions—Mountain and Pacific. The effect of age is negative for small firms (the older the firm, the fewer durable jobs it generates), but it's positive for large firms. Finally, being part of a multi-plant firm enhances the number of durable jobs a firm creates, although the effect tends to be weaker for larger firms.

In summary, this article presents evidence on the turnover of firms in U.S. manufacturing and estimates of the durability of jobs generated by different firms. A primary conclusion arising from the estimates is that a job generated at a large firm tends to last longer than a job generated at a small one.

REFERENCES

Davis, Steven, John Haltiwanger, and Scott Schuh. "Published Versus Sample Statistics from the ASM: Implications for the LRD," Proceedings of the American Statistical Association, Business and Economics Statistics Section (1990).

Davis, Steven, and John Haltiwanger. "Gross Job Creation and Destruction: Microeconomic Evidence and Macroeconomic Implications," *NBER Macroeconomic Annual*, 5 (1990), pp. 123-68.

Davis, Steven, and John Haltiwanger. "Gross Job Creation, Gross Job Destruction, and Employment Reallocation," *Quarterly Journal of Economics,* 107 (1992), pp. 819-63.

Demsetz, H. "Industry Structure, Rivalry and Public Policy," *Journal of Law and Economics,* 16 (1973).

Dixit, A. "Investment and Hysteresis," *Journal of Economic Perspectives,* 6 (1992), pp. 107-32.

Dixit, A., and Rob, R. "Switching Costs and Sectoral Adjustments in General Equilibrium with Uninsured Risk," *Journal of Economic Theory* (forthcoming).

Dunne, Timothy, Mark Roberts, and Larry Samuelson. "Patterns of Firm Entry and Exit in U.S. Manufacturing Industries," *Rand Journal of Economics,* 19 (1988), pp. 495-515.

Dunne, Timothy, Mark Roberts, and Larry Samuelson. "Plant Turnover and Gross Employment Flows in the U.S. Manufacturing Sector," *Journal of Labor Economics,* 7 (1989a), pp. 48-71.

Dunne, Timothy, Mark Roberts, and Larry Samuelson. "Firm Entry and Postentry Performance in the U.S. Chemical Industries," *Journal of Law and Economics,* 32(2), Part 2, 1989(b), pp. S233-71.

Dunne, Timothy, Mark Roberts, and Larry Samuelson. "The Growth and Failure of U.S. Manufacturing Plants," *Quarterly Journal of Economics,* 104 (1989c), pp. 671-98.

Ericson, Richard, and Ariel Pakes. "An Alternative Theory of Firm and Industry Dynamics," Working Paper, Yale University (1989).

Evans, David. "Tests of Alternative Theories of Firm Growth," *Journal of Political Economy,* 95 (1987a), pp. 657-74.

Evans, David. "The Relationship Between Firm Growth, Size and Age: Estimates for 100 Manufacturing Industries," *Journal of Industrial Economics,* 35 (1987b), pp. 567-81.

Hall, B. "The Relationship Between Firm Growth, Size and Age: Estimates for 100 Manufacturing Industries," *Journal of Industrial Economics,* 35 (1987), pp. 583-606.

Hart, P.E., and S.J. Prais. "The Analysis of Business Concentration," *Journal of Royal Statistical Society,* 119 (part 2, 1956), pp. 150-81.

Hopenhayn, H. "Entry, Exit, and Firm Dynamics in Long-Run Equilibrium," *Econometrica,* 60 (1992a), pp. 1127-50.

Hopenhayn, H. "Exit, Selection and the Value of Firms," *Journal of Economic Dynamics and Control,* 16 (1992b), pp. 621-53.

Jovanovic, B. "Selection and the Evolution of Industry," *Econometrica,* 50 (1982), pp. 649-70.

Lambson, Val. "Industry Evolution with Sunk Costs and Uncertain Market Conditions," *International Journal of Industrial Organization,* 9 (1991), pp. 171-96.

Lambson, Val. "Competitive Profits in the Long-Run," *Review of Economic Studies,* 59 (1977), pp. 125-42.

Leonard, J. "In the Wrong Place at the Wrong Time: The Extent of Frictional and Structural Unemployment," in D. Lang and J. Leonard, eds., *Unemployment and the Structure of Labor Markets.* New York: Basil Blackwell, 1987.

Lippman, S.A., and R. P. Rummelt. "Uncertain Imitability: An Analysis of Interfirm Differences in Efficiency Under Competition," *Bell Journal of Economics,* 13 (1982), pp. 418-38.

Lach, S. and R. Rob. "R&D, Investment, and Industry Dynamics," University of Pennsylvania CARESS Working Paper 92-02.

Lucas, R., and E.C. Prescott. "Investment Under Uncertainty," *Econometrica,* 39 (1971), pp. 659-81.

Orr, D. "The Determinants of Entry: A Study of the Canadian Manufacturing Industries," *Review of Economic Statistics,* 56 (1974), pp. 58-66.

Pakes, A., and R. Ericson. "Empirical Implications of Alternative Models of Firm Dynamics," Working Paper, Yale University, 1990.

Pindyck, R.S. "Irreversibility, Uncertainty and Investment," *Journal of Economic Literature,* 26 (1991), pp. 1110-48.

Rob, R. "Learning and Capacity Expansion Under Demand Uncertainty," *Review of Economic Studies,* 58 (1991), pp. 655-75.

Simon, H.A., and C. Bonini. "The Size Distribution of Business Firms," *American Economic Review,* 48 (1958), pp. 607-17.

Troske, K. "The Dynamic Adjustment Process of Firm Entry and Exit in Manufacturing, Finance, Insurance, and Real Estate," U.S. Census Bureau Working Paper (March 1993).

Troske, K. "The Importance of the Entry-Exit Process in Accounting for Industry Change," U.S. Census Bureau Working Paper (February 1993).

APPENDIX

Estimating the Number of Discounted Job Years Generated by Different-Sized Firms

Given the data in Table 2, the first step is to estimate the probabilities of increase and decrease in the size of firms as a function of their present size. It is assumed that firms can only switch to adjacent size groups and that the probability of switches are given by:

$$P_{i,i-1}=NEG_i[1+(i-1)k_1]/2,$$

$$P_{i,i+1}=POS_i[1+(i-1)k_2]/2,$$

$$P_{i,i}=1-P_{i,i-1}-P_{i,i+1}$$

where POS_i is the actual rate of growth of firms in size class i (i=1,...,5), NEG_i is the actual rate of decline, and k_1 and k_2 are parameters to be estimated. This generates a Markov chain, and we can compute its steady state. The numerical values that k_1 and k_2 take are chosen so that the steady state is as close as possible (in a chi-square sense) to the actual size distribution of firms (as given by the last row of Table 2).

Given this transition matrix, we can determine the expected number of jobs that each firm can generate during the course of its existence. For firms in class i this number is denoted n_i. Let $\delta=1/1.04$ represent the discount factor used to add up jobs generated in different periods, and let z_i represent the number of (current) jobs for firms in class i. For firms in class 1, which corresponds to firms with 1 to 99 employees, $z_1=50$; in general, z_i is the midpoint of the class i interval (and for firms in the 1000+ size class, $z_i = 1500$). Then the following system of equations determines n_i:

$$n_i=z_i+\delta(P_{i,i-1}n_{i-1}+P_{i,i}n_i+P_{i,i+1}n_{i+1}),\ i=0,1,\ldots,5, \text{ where } n_{-1}=0.$$

The solution to this system (the numerical values obtained for n_i's) is in the first row of Table 3.

To generate the rest of Table 3 we repeat the same simulations, except that POS_i and NEG_i are first adjusted to reflect the region-specific rate of growth/decline. A similar procedure is followed for age effects (Table 4) and for single- vs. multiple-plant firms (Table 5).

Making Money in the Housing Market: Is There a Sure-Fire System?

*Theodore M. Crone**

Almost two-thirds of the nation's households own the house in which they live. Although more than half of those houses are mortgaged, homeowner equity constitutes about one-third of all household wealth in the United States. For most people the largest single investment they will ever make is their house. It's not surprising, then, that housing market conditions command a great deal of attention in the media and that economists have devoted considerable effort to examining how well housing markets work.

Prospective buyers shop for houses both as consumers and as investors. As consumers, they are most interested in the characteristics of the house and the neighborhood—the number of bedrooms and baths, the presence of a garage or central air conditioning, and the quality of the schools to which they will send their children. As investors, buyers are interested in the return they are likely to realize on the house when it comes time to sell. The investment aspect of purchasing a house may receive little attention in discussions between prospective buyers and real estate agents, but it is an important concern for the buyer. In surveys of four different housing markets, Karl Case and Robert Shiller (1988) found that, for 44 to 64 percent of buyers, the investment factor was a major consideration in their decision to buy. And in only one of the markets were more than 5 percent of the houses bought

*Ted Crone is assistant vice president in charge of the Regional Economics section in the Philadelphia Fed's Research Department.

as rental properties. Less than 10 percent of buyers professed that the investment aspect was not a consideration at all. Prospective buyers, then, should value any information that could help in predicting the return they will realize on their investment. If there is some way to identify houses likely to yield unusually high returns, homeowners or investors could potentially profit by buying and selling simply to reap the better-than-average returns. If someone could exploit publicly available information to earn abnormal returns in housing, the housing market would not be efficient.

THE NORMAL RETURN TO HOUSING

Like the return to any other asset, the return to housing equals the cash flow (actual or imputed) from the asset plus any capital gain or loss, i.e., any increase or decrease in the value of the property.[1] In the case of stocks, the cash flow is simply the dividends earned. In the case of rental housing, the cash flow to the landlord is equal to the rent received minus the expenses incurred. On an annual basis, the pre-tax *rate of return* on rental housing equals the yearly rent minus operating and maintenance costs plus capital gain, all divided by the value of the house at the beginning of the year. Economists think about the return to owner-occupied housing in a similar way, except that there are no actual cash flows for rent or for any repairs that the homeowner performs himself. So, for homeowners, economists impute cash flows for rent and maintenance equal to what they would receive or spend for equivalent rental properties.

Tax considerations further complicate the calculation of the rate of return on rental and owner-occupied housing. A landlord is allowed to depreciate the property for tax purposes and to deduct the cost of maintenance, but he must also pay tax on any capital gains that he has realized on the property when he sells it.[2] Homeowners do not get to deduct housing depreciation on their income taxes, but they do not pay tax on the imputed rent they receive. Moreover, most homeowners are exempt from tax on the capital gains they receive from their primary residence.[3]

In calculating the rate of return on housing the capital gain is difficult to determine for houses that are not sold. For stocks held in one's portfolio, the prices of identical stocks sold in the market provide a precise measure of the capital gain. But houses are seldom identi-

---

[1]To simplify this discussion we will concentrate on the return on assets rather than the return on equity. We will assume that houses are not mortgaged and the equity of the homeowner or landlord is equal to the value of the asset. For those landlords or homeowners whose houses are mortgaged, their own equity in the property is less than the value of the house, and the mortgage interest payments are part of the annual costs.

[2]Algebraically the *after-tax* rate of return on rental property equals

$$\frac{(1-\tau_i)(R_t-M_t-D_t) + (1-\frac{\tau_g}{(1+r)^{S-t}})(A_t+D_t)}{V_t}$$

where R_t equals rent in period t, M_t equals maintenance costs in period t, D_t equals the depreciation allowance in period t, A_t equals the change in the market value in period t, V_t equals the value of the house at the beginning of the period, τ_i equals the income tax rate, τ_g equals the capital gains tax rate, S equals the period when the house is sold and capital gains taxes are paid, and r equals the discount rate.

[3]Thus, the after-tax return for homeowners is much simpler than that for landlords. It equals

$$\frac{R_t-M_t+A_t}{V_t}$$

where R_t is the imputed rent the homeowner receives in period t, M_t is the maintenance cost including the value of any labor on maintenance by the homeowner in period t, A_t is the change in the market value in period t, and V_t is the value of the house at the beginning of period t.

cal, and they are sold at infrequent intervals. Therefore, analysts regularly use some average measure of housing price increases or decreases in the local market to estimate the capital gain or loss for houses not for sale.[4]

The problems in measuring the rate of return on housing in general and on owner-occupied housing in particular make it difficult to determine the long-run average rate of return to residential real estate. But according to the best estimates, the annual compound rate of return on residential real estate from the late 1940s to the early or mid-1980s was between 7.4 percent and 8.1 percent. The rate of return on stocks over that period was considerably higher—11 percent or more. On the other hand, the rate of return on Treasury bills was lower—less than 5 percent per year.[5] Why should the rate of return on housing differ from the return on other assets? One reason is that the risks for holding different types of assets are not the same. Investors demand a higher average return on those assets that entail more risk. An asset's "normal rate of return" includes the return on a risk-free asset, such as a Treasury bill, plus a risk premium, that is, an additional amount to compensate the investor for the added risk.

A common measure of the riskiness of an asset is how widely its return varies over time. Consider two stocks whose returns move up and down together, one in a range of -2 percent to 6 percent and the other in the range of -5 percent to 9 percent. The first is less risky than the second, and investors would demand a higher average return on the second stock to compensate for the additional risk or uncertainty about the return they will receive. The variation in the annual return to housing has been considerably lower than the variation in stock returns.[6] Consequently, the average long-run rate of return on housing has been lower than the return on stocks. The opposite has been true of the relative return on housing and Treasury bills. Both the variation in the rate of return and the average return have been higher for housing than for Treasury bills, reflecting the fact that housing is a riskier investment than Treasury bills.

Not only does the variation in the return to residential real estate over time differ from the variation in the returns to other assets, but in any one period there is a great deal of variation within the housing market itself. The return on housing varies from market to market and even among houses in the same local market.[7] Of course, the rates of return on different stocks also vary because of the differing fortunes of one company versus another. A stock investor, however, can protect herself against the unforeseen bad fortune of a particular company by diversifying her portfolio, that is, by buying a number of different stocks that reflect the overall market or by buying shares in a mutual fund that diversifies its holdings. Unexpected bad fortune for one company rep-

[4]A sizable literature has been developed on estimating appreciation rates for single-family houses. See Crone and Voith and the special issue of *The Journal of the American Real Estate and Urban Economics Association*, Vol. 19, No. 3 (Fall 1991).

[5]See Ibbotson and Fall, Ibbotson and Siegel, and Goetzmann and Ibbotson.

[6]See Ibbotson and Fall, Ibbotson and Siegel, and Goetzmann and Ibbotson. Ibbotson and Siegel estimated that between 1947 and 1982 the standard deviation (a common measure of variability) of annual returns to a portfolio of stocks on the New York and American exchanges and in the over-the-counter market was four and a half times as great as the standard deviation of the annual return to residential real estate. Goetzmann and Ibbotson estimated that the standard deviation of returns on the stocks in the S&P 500 was more than three times the standard deviation of the return on housing between 1947 and 1986.

[7]For evidence on the variation in returns between markets see Case and Shiller (1987), and for variations within markets, see Kiel and Carson.

resented in the portfolio is likely to be matched by unexpected good fortune for another company in the portfolio. This kind of diversification in the housing market is not available to the ordinary homeowner. Her entire real estate investment is likely to be in one house and subject to the fortunes of a single, local market. There is no mechanism for the homeowner to distribute her equity over a large number of houses in different markets. Theory suggests that this inability of homeowners to diversify their investment raises the rate of return buyers must expect before they are willing to invest in a house.[8]

Besides the inability to diversify their investment in housing, homeowners face another problem that does not confront the holders of stocks and bonds. It is not easy to sell a house quickly; or in the jargon of financial markets, houses are not very liquid. Unlike stocks and bonds that are traded frequently in organized markets with large numbers of buyers and sellers, houses are sold in markets where bids are received rather infrequently and the final price is usually the result of some negotiation between the buyer and the seller. If the seller needs the equity in a house quickly, she may have to sell at a price lower than one she might have negotiated under normal circumstances. Theory suggests this lack of liquidity for housing would tend to raise the normal return expected by a buyer.

Because of the special features of the housing market it is not possible to identify an investment that corresponds exactly to housing in terms of risk. Patric Hendershott and Sheng-Cheng Hu have suggested that the closest alternative to an investment in owner-occupied housing is a portfolio of mortgages. Since some of the features of owner-occupied housing, such as the inability to diversify or the lack of liquidity, do not apply to the mortgage market, housing and mortgages are not exact substitutes in terms of risk. Nevertheless, a comparison of the return to housing with mortgage returns shows periods when housing has earned a higher return than mortgages and periods when it has earned a lower return.

Can the savvy investor predict when housing will earn a higher-than-normal rate of return?

Hendershott and Hu compared the return on housing to the after-tax return on mortgages for various overlapping eight-year periods between 1956 and 1979. From 1956 to 1963 a homeowner in the 30 percent tax bracket would have earned 6.53 percent *less* per year after taxes by investing in an owner-occupied home than by investing in mortgages. From 1960 to 1967 the homeowner would have made .05 percent more per year by investing in a house than in mortgages. From 1968 to 1975 the homeowner would have made 7.1 percent *more* per year by investing in a house than by investing in mortgages.

The return on housing relative to other assets has clearly fluctuated over time. Moreover, in any one period, returns in some housing markets are clearly higher than returns in others. For example, the appreciation rate for housing in the Boston metropolitan area from 1983 through 1985 was 20 percent a year or higher, while in Los Angeles it was less than 7 percent in each of those years. On the other hand, from 1987 through 1989 Los Angeles had appreciation rates ranging from 11.5 percent to 27.9 percent while Boston's rates were 6.2

[8]Case, Shiller, and Weiss have suggested the creation of index-based futures and options markets to offer the homeowner protection against the risk associated with this inability to diversify.

percent or lower.[9] If the relative risks associated with housing in general or with specific housing markets change over time, these fluctuations in returns can be explained as changes in the risk premium investors demand for investing in housing.[10] But if, as most studies assume, the relative risks do not change, these fluctuations indicate that there have been periods of abnormal returns to housing. But can the savvy investor predict when housing in general or certain housing markets will earn a higher-than-normal rate of return?

Most discussions of market strategy, whether by economists or by financial advisors, have focused on the stock and bond markets. For more than 20 years economists have debated whether it is possible to systematically "beat the market," that is, to earn profits above those earned on similarly risky assets by consistently predicting abnormal returns. *Private or insider information* can be used to earn abnormally high returns on stocks.[11] The real question is whether *publicly available information* can be used in the same way. If the current price of a stock fully reflects all publicly available information, the information cannot be used to earn a higher-than-normal return and the market for the stock is said to be *efficient*.[12] Any new information relevant to future earnings is immediately reflected in the price of the stock. Housing markets are different from stock markets, so any conclusions about the efficiency of the stock market do not necessarily apply to the housing market. But the basic questions about the market's efficiency remain the same. Is all publicly available information reflected in housing prices, or can investors systematically make an abnormally high return from this information?

PREDICTING THE RATE OF RETURN ON HOUSING

Before a prospective buyer can make a profit in the housing market from publicly available information, that information must help him forecast future returns. Among the available public information, a buyer might consider past rates of return or appreciation, or population, job, and income growth to forecast the return on his investment. The simplest forecast models in terms of collecting data use only past returns to predict future returns or past appreciation to predict future appreciation. Therefore, researchers draw a distinction between using only past returns or appreciation and using all publicly available information to forecast future returns or appreciation.[13]

What Can Past Rates of Return and Appreciation Tell Us? One clue to the future return on any asset is past returns. For example, most estimates of the normal return to stocks, bonds,

---

[9]See Case and Shiller (1994).

[10]Unfortunately, Hendershott and Hu do not provide any statistics on the variation in mortgage rates or in the return to housing during the periods they examine. Such a statistic would indicate whether the relative risks between the two assets had changed. Nor do Case and Shiller (1994) give any measure of relative risks in the Los Angeles and Boston housing markets in the 1980s.

[11]See the article by Jaffe.

[12]For an early article on the efficient markets hypothesis see Fama. For a review of the literature see LeRoy. Certainly, market analysts are rewarded for identifying stocks that are undervalued, but any above-average return they may receive from trading these stocks should simply cover the costs of obtaining the necessary information and performing the analysis plus a normal return on the investment. As these informed analysts purchase the stocks and bid up the prices, those who have not borne the costs of obtaining the information and doing the analysis will not share in any extra profits derived from this information. See Grossman and Stiglitz.

[13]If one cannot forecast abnormal returns using only past rates of return or appreciation, the market is said to be *weak-form efficient*. If one cannot forecast abnormal returns using any publicly available information, the market is said to be *semi-strong efficient*.

or housing are based on the long-run average return for these assets. These normal returns are generally expressed in real terms, that is, after taking account of inflation. But knowing the long-run average or normal return to housing is of little help to the home buyer. The opportunity to make a better-than-average profit in the housing market depends on one's ability to forecast above-normal returns from past returns. Depending on the historical pattern of returns, an investor in the housing market might adopt one of two strategies. If local housing markets with abnormally high returns in one year generally experience abnormally high returns in succeeding years, the investor would purchase a house in a market where returns had been abnormally high on the assumption that these higher returns would continue. On the other hand, if abnormally low returns have historically been followed by abnormally high returns, the investor would buy a house in a market that had just experienced relatively low returns. If either strategy worked consistently to produce abnormally high returns for the investor, the housing market could be said to be inefficient because prices do not incorporate all the information available to the buyer and seller.

For homeowners or investors, a higher-than-normal return could come in the form of a higher rent (actual or implicit) or higher-than-average appreciation or both. Some researchers have estimated total returns to housing (rent minus operating costs plus capital gain), but more often researchers have concentrated on the capital gain component of the return, that is, the appreciation.

In a 1987 study, Karl Guntermann and Richard Smith used price data on FHA-financed houses in 57 metropolitan areas to compute a yearly appreciation rate for each market from 1968 to 1982. They also estimated what they consider to be the "normal" relationship between each metropolitan area's housing appreciation and the average for all 57 areas. Yearly appreciation was deemed abnormally high or low depending on how far it deviated from this estimated relationship. The authors then looked for a pattern of abnormally high or low appreciation that was offset after some fixed number of years. They found no significant correlation between abnormal appreciation in one year and abnormal appreciation in each of the succeeding five years, but they did find some offsetting appreciation rates in years 6, 7, 8, and 10.[14]

In contrast to Guntermann and Smith's results, more recent studies have found a positive correlation in abnormal returns over short periods of time. Using information on houses that were sold more than once between 1970 and 1986, Karl Case and Robert Shiller (1989 and 1990) estimated housing appreciation rates for four different metropolitan areas (Atlanta, Chicago, Dallas, and San Francisco). In two (Chicago and San Francisco) of the four areas, yearly appreciation rates were positively related to appreciation rates in the previous year.[15] Case and Shiller also used local rental

---

[14]While this pattern held in general, for some episodes abnormal appreciation was not offset but rather enhanced in years 7 and 8. Guntermann and Smith also examined patterns of abnormal appreciation after controlling for higher- or lower-than-average rental yields. With these adjustments there is no significant correlation of abnormal appreciation rates for years one through three, but there are significant correlations for years 4, 6, 7, and 8. Because of the statistical technique used to estimate this "normal" relationship between appreciation in each metro area and the national average, positive (negative) deviations from the relationship in one year will necessarily be offset by negative (positive) deviations in other years. There was no necessity, however, for the intervals between the offsetting appreciation rates to be the same across cities as Guntermann and Smith found.

[15]The authors also found a significant positive relationship when the data for all four metro areas were combined. For the San Francisco area, Case and Shiller used data only from Alameda County. In a similar study Dogan Tirtiroglu found that appreciation rates in selected communities in

indexes to develop a measure of excess returns and found that in three of the markets (Chicago, Dallas, and San Francisco), abnormally high returns in one year were positively correlated with abnormally high returns in the previous year.[16] James Poterba confirmed these results using data on 39 metropolitan areas from the National Association of Realtors. The findings of Richard Meese and Nancy Wallace also support the persistence of abnormal returns in the short run. Meese and Wallace found that during the 1970s and 1980s the rate of return to housing in 14 of 16 municipalities in the San Francisco metropolitan area was significantly related to rates of return in the previous three years.

Most studies have looked for higher-than-normal returns in local housing markets over relatively short periods of time, but Joseph Gyourko and Richard Voith examined relative appreciation rates over a longer span of time. For the period between 1971 and 1989 they identified only two markets (San Francisco and San Jose) from among 56 metropolitan areas that had significantly higher-than-average appreciation. Gyourko and Voith also found no consistent pattern of abnormal appreciation within individual markets. Abnormally high or low appreciation rates tended to persist for some time in five of the 56 markets. But in three of the markets abnormally high or low appreciation tended to be offset in the near term.

The evidence on whether future increases in house prices can be predicted by past increases alone is not conclusive. Even when Case and Shiller (1989) found evidence of a positive correlation between price increases in one year and increases in the following year, they also found that information on recent appreciation in the local market was not helpful in predicting the appreciation of individual houses.

Do Other Data Help? Past rates of return or housing appreciation, of course, are not the only information that might help in forecasting future returns or appreciation. There are good theoretical reasons to believe that other demographic and economic variables such as population growth, income growth, or construction costs could influence the appreciation rate or rate of return on single-family housing. Therefore, researchers have looked at the pattern of appreciation rates in combination with other variables.[17] Some of these other factors have often proven to have an independent effect on future appreciation, at least in the short run.

Increases in population, especially in the adult population, clearly increase the amount of housing demanded (see the study by Mills). But how do house prices respond to predictable changes in the adult population? In a widely quoted article, Gregory Mankiw and David Weil argued that house price appreciation in real terms is closely linked to the current growth in the population over 20 years of age. Since the size of this population group is predictable 20 years in advance, Mankiw and Weil concluded that housing price increases or decreases should be predictable many years into the future. But Americans who bought houses in the 1950s and 1960s apparently did not take the post-World War II baby boom into account; they did not bid up the price of housing in anticipation of "predictable" higher appreciation rates in the 1970s. Prices rose only as

the Hartford metropolitan area were positively correlated with the previous year's appreciation in neighboring communities.

[16]These results on total returns indicate that, at least for these housing markets, higher appreciation rates are not offset by reduced cash flow.

[17]They have almost always found that when these other factors are taken into account, recent past appreciation is significant in explaining current appreciation. See Hamilton and Schwab, Case and Shiller's 1990 study, Poterba, DiPasquale and Wheaton, and Abraham and Hendershott.

the baby boomers reached their adult years. Mankiw and Weil cite this episode as evidence that housing markets are not efficient.

But Stephen Holland has questioned the causal connection between the coming of age of the baby boomers and the rise in housing prices in the 1970s. His statistical tests showed that, over the long run, real house prices did not necessarily move together with the growth of housing demand. In another challenge to Mankiw and Weil, James Poterba examined how their measure of housing demand estimated from the size of the population adjusted for the age distribution affected real house prices in 39 metropolitan areas. When per capita income and construction costs were taken into account, no statistically significant relationship emerged between increases in the real price of houses and the demographically determined demand for housing. This is consistent with James Follain's earlier results that, in the long run, the cost of new housing net of land will be determined by the cost of supplying houses, that is, by construction costs.

Just because we can predict future price increases...does not necessarily imply that the market is not working well.

While the long-term relationship between population growth and housing appreciation continues to be debated, two studies have provided evidence that over the short term faster population growth does lead to higher appreciation. In several variations of their model of the housing market, Case and Shiller (1990) found that faster population growth in one year was related to higher appreciation or excess returns for housing in the following year. Bruce Hamilton and Robert Schwab came to the same conclusion in their study of house price appreciation in 49 metropolitan areas. Thus, the evidence for a short-term effect of population growth on house prices is stronger than the evidence for a long-term effect.

The demand for housing is fueled not only by population growth but also by income growth. As incomes rise, more and more individuals or families are able to set up separate households, and people are able to spend more money on housing. But how does this affect the appreciation of a typical house? One fairly consistent result in the literature on housing prices is that greater increases in real income or real income per capita in one year lead to greater increases in real house prices in the following year. Studies by Hamilton and Schwab and Poterba directly support this conclusion. Although Case and Shiller (1990) found little evidence that an increase in real income resulted in a near-term increase in real house prices, they did find evidence that, taking the implicit rent into consideration, real income growth did increase excess returns to housing. A recent study by John Clapp and Carmelo Giaccotto provided indirect evidence of the effect of income growth on house price appreciation. They found that a decline in the unemployment rate in one year was associated with an increase in real house prices in the following year. And usually a decline in the unemployment rate represents an increase in real income.[18]

The weight of the evidence suggests that increases in population and income can result in higher real house prices at least over a one- to three-year period. But the price of housing is

[18]Clapp and Giaccotto did not control for changes in real income. Case and Shiller (1990) found that once income growth was taken into account, employment growth did not affect future price increases or excess returns.

determined not only by factors that affect demand but also by supply considerations such as the cost of building new houses. Are construction costs, then, an indicator of future changes in house prices? Case and Shiller (1990) addressed this question in their study of four major metropolitan areas, and the results were mixed. In many cases, they found that the higher the ratio of construction costs to price, the higher were housing price increases or excess returns in the following year. But this result depended on which metropolitan area was being considered and which other factors were being taken into account.

In trying to identify information that would help predict future increases in house prices, most researchers have looked at the fundamental factors driving housing demand (population and income) or the cost of supplying housing (construction costs). Peter Linneman, however, has taken another tack. Using data from the *Annual Housing Surveys* for the Philadelphia metropolitan area in 1975 and 1978, he identified houses that were undervalued in 1975 based on what their characteristics (number of bedrooms, central air conditioning, etc.) suggested the house should be worth.[19] He found that houses that were undervalued relative to their characteristics in 1975 appreciated more than other houses in the following three years. In many ways, Linneman's experiment mirrors what happens in the housing market. Home buyers shop for the best value based on the features of the house. They purchase the one that has the features they want at the lowest price. Linneman admits that the higher appreciation he observed was not enough to offset any transactions costs that a short-term investor would incur if he tried to buy these undervalued houses and sell them within a three-year period.[20] Thus, the short-term investor could not profit from such a strategy.

Various attempts over the past decade to find indicators of future appreciation of house prices have been at least partially successful. While there is little evidence of any ability to predict abnormal appreciation over the long term, a number of studies have identified indicators of abnormal appreciation over periods ranging from one to three years.

BUT WHO CAN PROFIT FROM THESE PREDICTIONS?

Just because we can predict future price increases from publicly available information does not necessarily imply that the market is not working well. We must also be able to systematically earn an above-normal return from these predictions before we can conclude that the market is not efficient. In February of each year, for example, a gas station operator may be able to predict that the price of gasoline will rise by the Fourth of July. He is not likely to make any abnormal profit from this information, however. If he buys gas for delivery in July, he will have to pay the higher price. If he buys extra gas in February to sell in July, the storage and carrying costs are likely to eat up any extra profit he would have made.

In the housing market, those who already own their homes will profit from any abnormal appreciation whether it is predictable or not. If they can actually predict a higher than normal appreciation over the near term, some may even delay selling their homes to realize

[19]Linneman used the common hedonic regression technique to determine the value of various housing characteristics in 1975. Unfortunately, Linneman did not have market prices for the houses in his sample but only the owners' estimates. He tried to overcome this limitation by redoing his experiment only with houses recently purchased in 1975 on the assumption that estimates by owners of these recently purchased houses would be close to the purchase price. He got the same results with this smaller sample.

[20]Short-term investors who purchase a property based solely on expected capital gain are sometimes referred to as speculators.

that return. But what about those who do not own a house in a market where higher than normal appreciation is predicted over the short term? For these potential buyers, certain features of the housing market make it difficult to earn abnormal returns.

A major difference between buying and selling a house and buying and selling financial investments such as stocks or bonds is the cost associated with the transaction. Discount brokers often charge 0.7 percent or less of the value of the stock to execute a purchase or sell order. For residential real estate the transaction costs include transfer taxes, deed-recording fees, title insurance, loan origination fees, and real estate commissions, and some of these costs can be substantial. For example, loan origination fees are typically 2 to 3 percent of the value of the mortgage, title insurance is usually 0.5 to 1.0 percent of the purchase price of the house, and real estate commissions are typically 6 percent or more of the selling price. For the investor in the housing market these transaction costs may be more of a hurdle than for the homeowner because the investor's after-tax return is likely to be lower. U.S. tax law favors homeowners over investors. Both get to deduct the interest on any mortgage on the property. But homeowners do not have to pay income tax on the implicit rent they receive, and in most cases the capital gain on their primary residence is also exempt from taxes. Therefore, a situation that may present an abnormal after-tax return to a homeowner may not present an abnormal return to an investor. Even for homeowners, however, the prospect of high short-term appreciation may not be enough to induce them to buy a house that is relatively far from where they work or different from the one they prefer. They will want to move to the house of their choice after the period of abnormal appreciation, but the cost of buying and selling a house over a short span of time as well as the cost of moving is likely to wipe out any excess profits the homeowner might have expected.

Besides the costs involved in actually buying and selling houses, the costs of gathering information in the housing market tend to discourage speculation. Housing markets are very localized even within metropolitan areas, and information about one market may not apply to other nearby markets. Information about what is for sale and recent sale prices is available for local markets. However, the prices of otherwise identical houses can differ greatly from one locality to another, and in this sense housing markets are local markets. The cost of gathering information for a small housing market may be almost as great as the cost for a large market, but the number of opportunities to profit from the information on a small market is limited. The investor has to weigh the cost of gathering the information against the profit he can expect to reap from the information. Prospective buyers already in the market for a house might increase their return by shopping around for an undervalued house. Whether they actually increase their return will depend on the costs of searching for the undervalued house, and these costs may vary from buyer to buyer. For example, a person who does not have to travel far to the neighborhood in which he intends to buy will incur lower costs than someone who must travel some distance to search for a house. Search costs may also be lower for someone who can take his time searching because he does not have to move quickly.

CONCLUSION

Most home buyers undoubtedly hope to make a better-than-average return on their investment, and there are even people who speculate in housing, buying units in neighborhoods where they anticipate higher-than-average appreciation. Indeed, some of them may realize abnormally high profits relative to the risk they take on certain investments. Others, of course, lose money on their investment.

But there is no convincing evidence that over the long term speculating in real estate produces abnormal profits relative to the risks involved. Recent economic research has indicated that there are some good indicators of higher-than-average appreciation rates *over the short term.* But the high cost of obtaining that information and of buying and selling houses suggests that investors may not be able to systematically make abnormally high profits in the housing market.

Does this mean that any prospective home buyer wastes his time looking for the "undervalued" house? Not necessarily. The literature suggests he may find such a house. Whether he earns higher-than-average profits depends on the costs of his search.

REFERENCES

Abraham, Jesse M., and Patric H. Hendershott. "Bubbles in Metropolitan Housing Markets," NBER Working Paper No. 4774 (June 1994).

Case, Karl E., and Robert J. Shiller. "Prices of Single-Family Homes Since 1970: New Indexes for Four Cities," *New England Economic Review* (September/October 1987), pp. 45-56.

Case, Karl E., and Robert J. Shiller. "The Behavior of Home Buyers in Boom and Post-Boom Markets," *New England Economic Review* (November/December 1988), pp. 29-46.

Case, Karl E., and Robert J. Shiller. "The Efficiency of the Market for Single-Family Homes," *American Economic Review,* 79 (March 1989), pp. 125-37.

Case, Karl E., and Robert J. Shiller. "Forecasting Prices and Excess Returns in the Housing Market," *Journal of the American Real Estate and Urban Economics Association,* 18 (Fall 1990), pp. 253-73.

Case, Karl E., and Robert J. Shiller. "A Decade of Boom and Bust in the Prices of Single-Family Homes: Boston and Los Angeles 1983 to 1993," *New England Economic Review* (March/April 1994), pp. 40-54.

Case, Karl E., Robert J. Shiller, and Allan N. Weiss. "Index Based Futures and Options Markets in Real Estate," *The Journal of Portfolio Management* (Winter 1993), pp. 83-92.

Chinloy, Peter. "Returns to Holding Housing," *Journal of Housing Economics,* 2 (1992), pp. 310-23.

Clapp, John M., and Carmelo Giaccotto. "The Influence of Economic Variables on Local House Price Dynamics," *Journal of Urban Economics,* 36 (1994), pp. 161-83.

Crone, Theodore M., and Richard P. Voith. "Estimating House Price Appreciation: A Comparison of Methods," *Journal of Housing Economics,* 2 (1992), pp. 324-38.

DiPasquale, Denise, and William C. Wheaton. "Housing Market Dynamics and the Future of Housing Prices," *Journal of Urban Economics,* 35 (1994), pp. 1-27.

Fama, E. F. "Efficient Capital Markets: A Review of Theory and Empirical Work," *Journal of Finance,* 25 (1970), pp. 383-417.

Follain, James R. "The Price Elasticity of the Long-Run Supply of New Housing Construction," *Land Economics,* 55 (1979), pp. 190-99.

Goetzmann, William N., and Roger G. Ibbotson. "The Performance of Real Estate as an Asset Class," *Journal of Applied Corporate Finance,* 3 (Spring 1990), pp. 65-76.

Grossman, Sanford J., and Joseph E. Stiglitz. "On the Impossibility of Informationally Efficient Markets," *American Economic Review,* 70 (June 1980), pp. 393-408.

Guntermann, Karl E., and Richard L. Smith. "Efficiency of Real Estate Markets," *Land Economics* 63 (1987), pp. 34-45.

Gyourko, Joseph, and Richard Voith. "Local Market and National Components in House Price Appreciation," *Journal of Urban Economics,* 32 (1992), pp. 52-69.

Hamilton, Bruce W., and Robert M. Schwab. "Expected Appreciation in Urban Housing Markets," *Journal of Urban Economics,* 18 (1985), pp. 103-18.

Hendershott, Patric H., and Sheng-Cheng Hu. "Inflation and Extraordinary Returns on Owner-Occupied Housing: Some Implications for Capital Allocation and Productivity Growth," *Journal of Macroeconomics,* 3 (1981), pp. 177-203.

Holland, A. Steven. "The Baby Boom and the Housing Market: Another Look at the Evidence," *Regional Science and Urban Economics,* 21 (1991), pp. 565-71.

Ibbotson, Roger G., and Carol L. Fall. "The United States Market Wealth Portfolio," *The Journal of Portfolio Management,* 6 (Fall 1979), pp. 82-92.

Ibbotson, Roger G., and Lawrence B. Siegel. "Real Estate Returns: A Comparison with Other Investments," *Journal of the American Real Estate and Urban Economics Association,* 12 (1984), pp. 219-42.

Jaffe, Jeffrey F. "The Effect of Regulation Changes on Insider Trading," *The Bell Journal of Economics and Management Science,* 5 (1974), pp. 93-121.

Kiel, Katherine A., and Richard T. Carson. "An Examination of Systematic Differences in the Appreciation of Individual Housing Units," *The Journal of Real Estate Research,* 5 (1990), pp. 301-18.

LeRoy, Stephen F. "Efficient Capital Markets and Martingales," *Journal of Economic Literature,* 27 (1989), pp. 1583-21.

Linneman, Peter. "An Empirical Test of the Efficiency of the Housing Market," *Journal of Urban Economics,* 20 (1986), pp. 140-54.

Mankiw, N. Gregory, and David N. Weil. "The Baby Boom, The Baby Bust, and the Housing Market," *Regional Science and Urban Economics* 19 (1989), pp. 235-58.

Meese, Richard, and Nancy Wallace. "Testing the Present Value Relation for Housing Prices: Should I Leave My House in San Francisco," *Journal of Urban Economics,* 35 (1994), pp. 245-66.

Mills, Leonard. "Understanding National and Regional Housing Trends," Federal Reserve Bank of Philadelphia *Business Review* (September/October 1991), pp. 15-23.

Poterba, James M. "House Price Dynamics: The Role of Tax Policy and Demography," in William C. Brainard and George L. Perry, eds., *Brookings Papers on Economic Activity 2,* Washington, DC, Brookings Institution, 1991, pp. 143-203.

Tirtiroglu, Dogan. "Efficiency in Housing Markets: Temporal and Spatial Dimensions," *Journal of Housing Economics,* 2 (1992), pp. 276-92.

FEDERAL
RESERVE BANK OF
PHILADELPHIA

Business Review Ten Independence Mall, Philadelphia, PA 19106-1574