Business Review

Federal Reserve Bank of Philadelphia

September • October 1991

ISSN 0007-7011

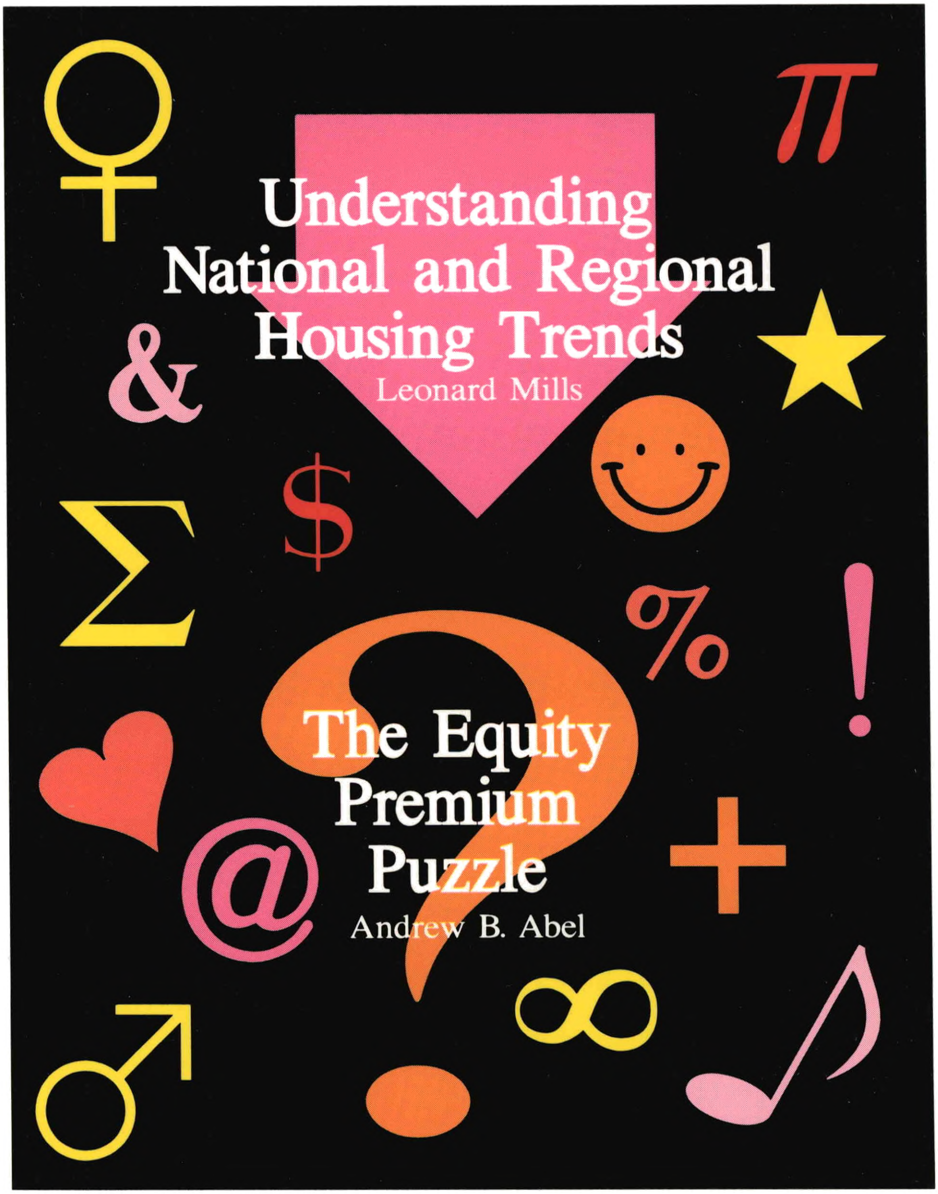

Business Review

The BUSINESS REVIEW is published by the Department of Research six times a year. It is edited by Patricia Egner. 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, 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.

SEPTEMBER/OCTOBER 1991

The Equity Premium Puzzle

*Andrew B. Abel**

The basic paradigm used by financial economists to explain rates of return on assets was called into question a few years ago by economists Rajnish Mehra of the University of California at Santa Barbara and Edward Prescott of the University of Minnesota. In a 1985 article published in the *Journal of Monetary Economics*, Mehra and Prescott presented a powerful argument that commonly used economic models were incapable of accounting for the historically observed rates of return on stocks and short-term bonds (bills). Specifically, they found that, in the 90 years from 1889 to 1978, the average real rate of return on stocks was 6.98 percent per year, while the average real rate of return on bills was only 0.80 percent per year. The rate of return on stocks minus the rate of return on bills—the so-called equity premium—averaged an astonishing 6.18 percent per year.

Why was the equity premium so large? The

*Andrew B. Abel is a Professor of Finance at the Wharton School, University of Pennsylvania, and a Visiting Scholar at the Federal Reserve Bank of Philadelphia. He thanks Stephen Cecchetti, Dean Croushore, Leonard Nakamura, Jeremy Siegel, Herb Taylor, and Stephen Zeldes for detailed comments, and Pat Egner for editorial assistance.

obvious answer is "risk." Stocks are much riskier than bills, and investors would not want to hold stocks unless they were compensated for the higher risk by earning a higher average rate of return. This basic insight—that investments with higher risk should earn higher average returns—underlies the capital asset pricing model (CAPM), initially developed in the 1960s and refined considerably in the last three decades.

Perhaps the most significant refinement, the consumption capital asset pricing model (CCAPM), recognizes that the ultimate reason for holding wealth is to provide for future consumption; as a result, the equity premium should depend on the variability of consumption and its relation to stock returns. In light of the small fluctuations in U.S. real consumption per capita, however, Mehra and Prescott found that the CCAPM could account for an equity premium of only 0.35 percent per year, a tiny fraction of the historically observed equity premium. To describe this large discrepancy, they coined the term "equity premium puzzle."

Trying to explain average rates of return over a historical time period is a much less formidable task than, say, trying to forecast the returns on stocks or bills in any particular year. Indeed, economists readily admit their limited ability to forecast asset returns. But the CCAPM's inability to account for average rates of return on stocks and bills, *even after the fact*, is a serious indictment of this model's practical value.

Moreover, the basic CCAPM is essentially the same as the model underlying the theory of long-run economic growth and the new strand of classical macroeconomics known as real-business-cycle theory. If the CCAPM has to be discarded or even drastically altered, then much of growth theory and new classical macroeconomics may need a major overhaul. Indeed, the equity premium puzzle could lead economists to reformulate basic models of decisionmaking in the presence of risk.

THE CONSUMPTION CAPITAL ASSET PRICING MODEL

The CCAPM is a sophisticated economic model of the prices and rates of return on assets. To understand its basic workings, let's first see how asset prices would be determined if investors did not care about the riskiness of their investments.

Risk-Neutral Investors. Confronted with two assets offering different expected rates of return, risk-neutral investors would buy the asset with the higher expected rate of return and sell the asset with the lower expected rate of return. These purchases and sales by investors, however, ultimately affect the expected rates of return. The asset with the higher expected rate of return would attract buyers, and its price would be bid upward. Of course, when the price of the asset increases, its rate of return falls because investors must pay more to receive its payoffs. Similarly, the asset with the lower expected rate of return would fall in price as investors sold it. The fall in price would increase the asset's expected rate of return by allowing investors to acquire ownership and future payoffs at a lower price.

The adjustment of asset prices and rates of return would cause the gap between the rates of return to shrink. When there is no more upward or downward pressure on asset prices, the asset markets are said to be in equilibrium, and the expected rates of return on both assets will be the same. Thus, with risk-neutral investors, the basic model of asset pricing predicts that asset prices will adjust until all assets offer equal expected rates of return.

Risk-Averse Investors. Most investors are anything but risk-neutral, demanding a higher expected rate of return in order to hold a riskier asset. But how do we measure the riskiness of an asset? The CCAPM offers a very precise answer. Instead of measuring the riskiness of an asset simply by the variability of its returns, the CCAPM uses the relationship between the asset's returns and the value an investor places on having an additional dollar of funds.[1] When the investor's overall wealth is low, his consumption is low and he places a relatively high value on an additional dollar of funds. And when the investor's overall wealth is relatively high, his consumption is relatively high and the value he places on an additional dollar of funds is relatively low.

According to the CCAPM, an asset is risky if its low payoffs occur when consumption is low (and the value of additional funds is high), and its high payoffs occur when consumption is high (and the value of additional funds is low). On the other hand, an asset would have negative risk if its high returns occur when consumption is low and its low returns occur when consumption is high; in this case, rather than being risky, the asset would provide insurance by offering high returns when the investor values additional funds most highly (when consumption is low).

The CCAPM predicts that risk-averse investors will choose assets with the *highest expected value of returns weighted by the value placed by investors on additional funds.* As in the case of risk-neutral investors, prices will adjust until equilibrium is reached. In equilibrium, the expected rates of return weighted by the value of additional funds will be the same for all assets.[2] Nevertheless, assets with relatively high risk will have higher average returns than assets with relatively low risk. The higher average return of a risky asset is offset by the fact that the high returns occur when additional funds have low value to investors.

If we apply the CCAPM to stocks and bills, the average rates of return weighted by investors' value of additional funds should be equal for stocks and bills. To the extent that stock returns (which comprise dividends plus capital gains or losses resulting from changes in the prices of stocks) are riskier than bill returns, the average rate of return on stocks should be higher than the average rate of return on bills. How much higher depends quantitatively on two factors: (1) the covariances of consumption growth with stock returns and bill returns, which measure the sizes of fluctuations in returns and how strongly these fluctuations are related to the fluctuations in consumption growth;[3] and (2) the coefficient of relative risk aversion, A, which indicates how much the value of additional funds increases when consumption falls.[4]

Mehra and Prescott combined a simple economic model conventionally used in growth theory and real-business-cycle theory with the actual historical variability of U.S. consumption to capture the covariances of consumption with asset returns. The value of A is an important ingredient in this analysis, and, based on their reading of theoretical and empirical research, Mehra and Prescott argued that conventionally accepted values for A lie between 0 and 10. Using a variety of values for A in this range, they found that, in the framework of the

[1]The value of additional funds is measured by what economists call "the marginal utility of consumption."

[2]Equilibrium is represented by the following technical condition: $E\{(1+r_1)*MU\} = E\{(1+r_2)*MU\}$, where MU is the marginal utility of consumption (the value of additional funds), r_1 and r_2 are the real rates of return on assets 1 and 2, respectively, and E { } denotes the expectation of the term that appears inside the brackets.

[3]Technically, the covariance of stock returns with consumption growth equals the product of the correlation coefficient between stock returns and consumption growth, the standard deviation of stock returns, and the standard deviation of consumption growth.

[4]If the coefficient of relative risk aversion equals A, then a 1 percent fall in consumption increases the value an investor places on an additional dollar of funds by A percent. For example, if A = 6, then a 2 percent fall in consumption increases the value of an additional dollar of funds by 12 percent.

CCAPM, they could not simultaneously account for an average equity premium higher than 0.35 percent per year and an average return on bills of less than 4 percent per year. For the average equity premium to be as large as the historically observed equity premium, the value of A would have to be extremely high, around 30 or 40, which is much higher than the conventionally accepted values for A.

To see why values of A around 30 or 40 are conventionally viewed as implausibly high, suppose that you face a risky situation that will either raise your total wealth by 50 percent or lower it by 50 percent, and each of these outcomes has a 50-50 chance of occurring.[5] How much would you be willing to pay for insurance to avoid this risky situation? If you were risk-neutral, so that A = 0, you would not care about risk and would pay zero for such insurance. However, if risk-averse, you would be willing to pay something for this insurance, and the amount would depend on the strength of your risk aversion measured by A.[6] With A = 2, you would be willing to pay 25 percent of your wealth; with A = 10, you would be willing to pay 46 percent; and with A = 30, you would be willing to pay 49 percent. Because it seems implausible that you would pay 49 percent of your wealth to avoid an even chance of losing 50 percent of your wealth or gaining 50 percent of your wealth, many economists reject as implausible values of A as high as 30.

REEXAMINATION OF THE DATA

One approach to reconciling the gap between the CCAPM and the average actual equity premium reported by Mehra and Prescott is to reexamine the historical data. The average rates of return on bills (0.80 percent per year) and stocks (6.98 percent per year) reported by Mehra and Prescott are based on 90 years of U.S. data. However, recent research by Jeremy Siegel (1991) indicates that the rates of return in the years between 1889 and 1978 may not have been truly representative of the underlying rates of return over a longer span of time. Siegel compiled annual rates of return on stocks and bills for the period from 1802 to 1990, starting 87 years before and ending 12 years after the period examined by Mehra and Prescott.[7] The variability of stock returns is much greater than the variability of bill returns, which is consistent with the notion that stocks are much riskier than bills.

Although the greater variability of stock returns is clear from Figure 1, the difference in the average rates of return on stocks and bills is not. To get a clearer view of the average rates of return, we can calculate the 30-year moving average rate of return, which, for any given year, is the average of the rates of return over the previous 30 years. In Figure 2, the difference between the 30-year moving averages of returns for stocks and for bills is the average of the equity premium over the previous 30 years. The 30-year moving average equity premium increased substantially during the 1940s and 1950s and remained high during the 1960s and 1970s.

The average rates of return calculated by Siegel for the period examined by Mehra and Prescott (1889-1978) differ somewhat from the values reported by Mehra and Prescott (see

[5]In addition to financial assets, total wealth includes all other tangible assets, such as real estate and consumer durables, and also human capital, which is the present value of a person's current and future labor income.

[6]The general formula is

$y = 1 - [(^1/2)(1 - x)^{1-A} + (^1/2)(1 + x)^{1-A}]^{1/(1-A)}$, where x is the fraction of your wealth that you could gain or lose with a 50-50 chance, A is the coefficient of relative risk aversion, and y is the fraction of your wealth that you would pay to avoid this risk.

[7]As in Mehra and Prescott, the average rates of return are arithmetic averages (rather than geometric averages) of annual rates of return.

FIGURE 1
Real Returns on Stocks and Bills
Annual returns 1802-1990

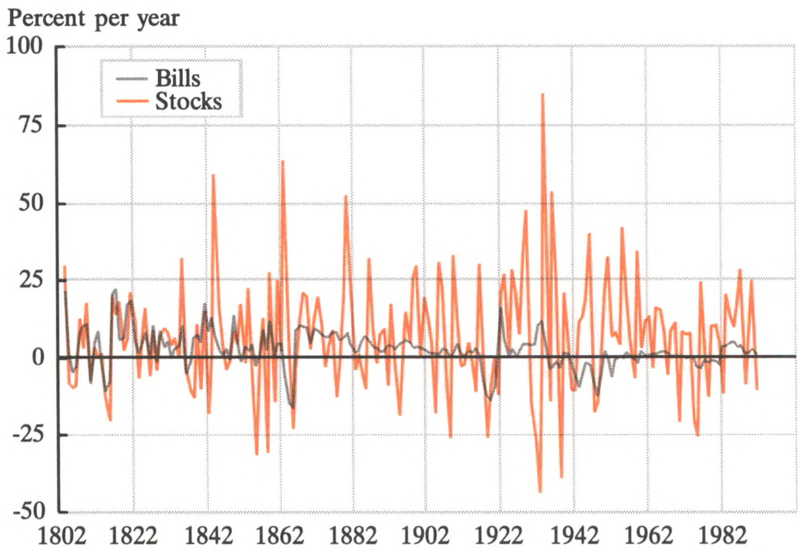

FIGURE 2
Real Returns on Stocks and Bills
30-year moving average, 1831-1990

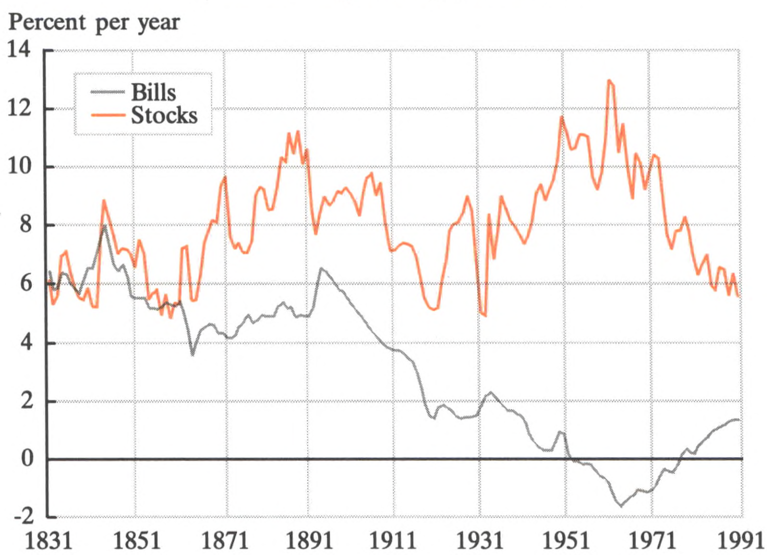

table). The differences arise because Siegel used a different stock price index, a different measure of inflation, and, for part of the period, a different short-term interest rate.[8] Despite these differences, the basic result is the same: the average equity premium from 1889 to 1978 was very large—well over 6 percent per year.

But including the additional 99 years of data in Siegel's study reduces the average equity premium from 6.96 percent per year to 4.62 percent per year. The reason for this drop is that the average real rate of return on bills rises to 3.19 percent per year when we include data over the entire 1802-1990 period; the average real rate of return on stocks is virtually the same over that period as over the period studied by Mehra and Prescott. However, even this lower value of the equity premium is much higher than that predicted by the CCAPM examined by Mehra and Prescott.[9]

Rates of Return and the Equity Premium

(Percent per year)

Period	Real Return on Bills	Real Return on Stocks	Equity Premium
1802-1888	5.62	7.52	1.90
1889-1978	0.91	7.87	6.96
1979-1990	2.73	9.44	6.71
1802-1990	3.19	7.81	4.62

Another way to examine the reliability of the historical average rates of return is to estimate how close the historical average rates of return are to the underlying rates of return investors expect when making their portfolio decisions. Applying statistical techniques to data from 1892 to 1988, Stephen Cecchetti, Pok-sang Lam, and Nelson Mark (1991) found that the average equity premium was 6.03 percent, but that the equity premium expected by investors could have been anywhere from 2.35 percent to 9.71 percent.[10] Even the low value of 2.35 percent for the equity premium is higher than the CCAPM studied by Mehra and Prescott can explain.

Because the equity premium still appears large after reexamining the historical data on returns, the next step is to reexamine the basic CCAPM.

EXTENSIONS OF THE BASIC CCAPM

The other approach to explaining the equity premium puzzle is to see if the basic CCAPM can be modified to produce a realistic value of

[8]The short-term real interest rate is intended to measure the short-term riskless rate of return, which is the real rate of return that can be earned on a short-term asset that has no risk of default or price variation. Siegel, as well as Mehra and Prescott, used the interest rate on short-term Treasury bills to measure the short-term riskless rate from 1920 onward. To measure the riskless interest rate before 1920, Mehra and Prescott used the short-term commercial paper rate, but Siegel adjusted the commercial paper rate to adjust for the risk of default by issuing companies.

[9]The predictions from the CCAPM studied by Mehra and Prescott are based on the variability of consumption growth during the period 1889-1978. Strictly speaking, we should use the variability of consumption growth during the period 1802-1990 to compare the predicted equity premium with the actual average equity premium reported by Siegel. However, there are no reliable annual data on consumption prior to 1889.

[10]More precisely, their statistical analysis indicates that if the expected equity premium was constant, then we can be 95 percent confident that it was in the range of 2.35 percent to 9.71 percent. As for the riskless rate, its average value was 1.15 percent, and we can be 95 percent confident that the expected value of the riskless rate was between -0.47 percent and 2.77 percent.

the average equity premium using a value of the coefficient of relative risk aversion, A, in the conventionally accepted range of 0 to 10. Several potential modifications are discussed below.

Richer Models of Underlying Risk. In their version of the CCAPM, Mehra and Prescott assumed that consumption fluctuations behaved according to a simple model that does not allow for the possibility of a large, sudden drop in consumption as might occur during a sharp depression. In addition, Mehra and Prescott assumed that fluctuations in stock dividends were matched exactly by fluctuations in consumption, and they used historical data on consumption to measure the variability of dividends.[11] Subsequent research, discussed below, has studied the importance of these assumptions by allowing for large, sudden drops in consumption and by allowing fluctuations in dividends to differ from fluctuations in consumption.

In a recent study, Thomas Reitz (1988) argued that if there is some possibility of a large, sudden drop in consumption accompanied by a large, sudden drop in dividends, then investors would be willing to hold stocks only if compensated by a high average equity premium. He found that extending the CCAPM to include the possibility of depressions with large, sudden drops in consumption could account for the historically observed equity premium. However, Mehra and Prescott (1988) point out that the potential depressions analyzed by Reitz involved declines in consumption of 25 percent or more during a single year. While it is true that consumption during the Great Depression fell 22 percent between 1929 and 1933,[12] Mehra and Prescott point out that in no single year did consumption fall as much as 9 percent.[13] Thus, they conclude that the drops in consumption in Reitz's study are too large to provide a realistic solution to the equity premium puzzle.

An alternative approach to modeling the riskiness of stocks is to incorporate in the model spans of good years (high consumption growth) and spans of bad years (low consumption growth), with unpredictable switches between the two. Shmuel Kandel and Robert Stambaugh (1990) and Cecchetti, Lam, and Mark (1991) used this approach, but concluded that a high value of A was still needed to explain the historically observed equity premium. Although this richer process of underlying risk did not help explain the average rates of return on stocks and bills, Kandel and Stambaugh point out that it helps explain other statistical features of returns, such as their predictability.

Another way to enrich the model of risk is to relax the assumption that fluctuations in dividends are matched exactly by fluctuations in consumption. One approach, followed by Cecchetti, Lam, and Mark (1991) and Kandel and Stambaugh (1990 and 1991), is to account for the fact that stocks are leveraged claims on firms. Firms generally raise capital by issuing both stocks and bonds. Because firms must pay their obligations to bondholders before they can pay dividends to stockholders, leverage tends to increase the riskiness of a stock and would increase the equity premium in the CCAPM. However, even taking account of historically observed degrees of leverage, a high value of A is still needed to account for the

[11]Dividends differ from stock returns because of changes in the price of stocks. The return on a stock equals the dividend plus the increase in the price of the stock (capital gain) or minus the decrease in the price of the stock (capital loss). In the CCAPM, the price of a stock is related to the current and future dividends weighted by the current and future marginal utilities of consumption. Given the behavior of consumption and dividends, we can compute the price of stock, and the rate of return on stock, using the CCAPM.

[12]Reitz (1988), footnote 9, p. 125.

[13]Mehra and Prescott (1988), p. 134.

historically observed value of the equity premium.

A more empirical approach to relaxing the assumption that fluctuations in dividends are matched exactly by fluctuations in consumption is simply to use historical data on dividends to measure dividend variability, and historical data on consumption to measure consumption variability. As pointed out by Cecchetti, Lam, and Mark (1991), dividends are much more variable than consumption.[14] Using the actual variability of dividends in the CCAPM raises the equity premium predicted by the CCAPM by about 50 percent for any given value of A.[15]

The general conclusion is that richer models of underlying risk can raise the value of the equity premium predicted by the CCAPM. However, the CCAPM still predicts a value for the equity premium that is much lower than the actual historical average value, if we continue to use a coefficient of relative risk aversion less than or equal to 10.

Differences Among Investors. The research discussed so far has assumed that investors are identical in all respects. Like other assumptions used in economic models, this one was made for the sake of simplicity. The question is whether this assumption is responsible for the small predicted value of the equity premium in most applications of the CCAPM.

To get an idea of the differences among investors and their portfolios, N. Gregory Mankiw and Stephen Zeldes (1991) studied the asset holdings of 2998 families. They found a striking degree of variation in the portfolios held. In particular, 72.4 percent of the families in the survey held no stocks at all.[16] Even among families that held more than $100,000 in other liquid assets, only 48 percent held stock. This finding is important because, to determine the prices of assets, the CCAPM typically uses the covariance of stock returns and *aggregate* consumption per capita.

But with almost three-fourths of the families holding no stock at all, the covariance should be calculated using the consumption not of *all* the families but only of those that hold stocks. Having made this change, Mankiw and Zeldes find that the covariance of stock returns and consumption per family *triples*, reflecting the facts that, compared to nonstockholders, stockholders have more volatile consumption and their consumption is more closely related to stock returns. This tripling of the covariance of stock returns and consumption reduces by about two-thirds the value of A needed to account for the equity premium. This finding is appealing, but leaves us asking why so many consumers—especially wealthy consumers

[14]In addition, the unpredictable components of dividend growth and consumption growth have a correlation coefficient of 0.443, which is lower than the value of 1.0 that is assumed in the Mehra-Prescott model.

[15]If the growth rates of consumption and dividends are jointly identically and independently distributed, the equity premium is approximately proportional to A times Cov(consumption growth, dividend growth). Using consumption growth to measure dividend growth in the CCAPM, the equity premium is approximately proportional to A times Var(consumption growth). Using data from Cecchetti, Lam, and Mark, Cov(consumption growth, dividend growth) = 0.002053 and Var(consumption growth) = 0.001398. Therefore, using actual dividend growth increases the equity premium by about 47 percent because 0.002053 is about 1.47 times as large as 0.001398.

[16]For the purposes of this study, a family that held stocks in a pension fund but did not directly own stocks was considered a nonstockholding family. Mankiw and Zeldes argue that this treatment is appropriate because only 49 percent of the labor force had a pension fund, and only 31 percent of these people had defined-contribution (rather than defined-benefit) plans. Thus, only 16 percent of the labor force had defined-contribution plans. In defined-benefit plans, the stocks held by the pension fund are more appropriately regarded as being owned by the employing firms rather than the worker because the firm bears the risk of changes in the value of stocks.

with large amounts of liquid wealth—hold no stock.

A more fundamental question is why consumption behaves so differently for different groups of consumers. The CCAPM is based on the assumption that even though individuals face idiosyncratic risks that do not hit everyone in the economy, they can protect their consumption from such risks by various sorts of risk-sharing and insurance arrangements. For example, life insurance, disability insurance, fire insurance, and so on protect an individual's consumption against various idiosyncratic risks. But problems such as the costs and difficulties of writing and enforcing various contracts prevent complete sharing of idiosyncratic risks. Theoretical studies have examined the impact of idiosyncratic risks on the equity premium,[17] but these studies do not provide empirical evidence of the importance of these factors in accounting for the equity premium puzzle. Further research in this area is needed.

Attitudes Toward Risk. Investors' attitudes toward risk are represented in economic models by utility functions that specify how much utility, or satisfaction, an investor gets for each possible level of consumption.[18] The most commonly used version of the CCAPM is based on a particular utility function with two important features: (1) consumption in any year affects utility in that year only; and (2) the utility function has a constant coefficient of relative risk aversion, which implies that the share of the portfolio held in risky assets does not depend on how much wealth the investor has. Utility functions with these features are convenient, but have an important limitation: they do not distinguish an investor's aversion to risk from his aversion to switching some consumption from one year to another year.

Kandel and Stambaugh (1991), Narayana Kocherlakota (1990), and Philippe Weil (1989) have investigated rates of return in the CCAPM using a more flexible utility function that distinguishes aversion to risk from aversion to substituting consumption between different years. However, they all conclude that, even with this more flexible structure, a very high value of the coefficient of relative risk aversion is needed to account for the historical value of the equity premium.

Moreover, Kandel and Stambaugh (1991) have suggested that the search for a version of the CCAPM that can explain a large equity premium with a value for A of less than 10 is perhaps misdirected. They argue that the conventional view that A is small (less than 10) is based on an unconvincing body of evidence. Furthermore, they point out that for risks that represent a relatively small portion of total wealth, high values of A may be plausible. For example, to avoid a risky situation that involves either a 1 percent gain or 1 percent loss of wealth with equal probabilities, a person with A = 30 would be willing to pay an insurance premium of 0.15 percent of his wealth (15 percent of the amount at risk), which is not implausible. Because high values of A (around 30) may be plausible for small risks, the important issue for asset pricing considerations is the degree of risk aversion appropriate for the magnitude of the risks investors bear in their portfolios. The value of A is extremely important for the equity premium puzzle because the CCAPM will produce a high value of the equity premium if A is large.

Until this issue is resolved, Kandel and Stambaugh urge us not to rule out high values of A, if we continue to use utility functions that

[17]See Mankiw (1986), Weil (1990), and Kahn (1988). A different aspect of differences among investors—different beliefs about future payoffs to risky assets—is examined in Abel (1989). That theoretical study shows that such differences tend to increase the equity premium predicted by the CCAPM.

[18]The marginal utility of consumption, discussed earlier, is the derivative of the utility function with respect to consumption.

have a constant coefficient of relative risk aversion. In light of the difference in plausible values of A for small and large risks, it may be appropriate to use more general utility functions for which the coefficient of relative risk aversion is not constant. Future research may pursue this suggestion.

Another modification of the attitude toward risk is to assume that an investor cares about his level of consumption relative to a benchmark or accustomed level of consumption attained in the recent past. So far, studies have taken two approaches to modeling an accustomed level of consumption. In one approach, dubbed "Catching up with the Joneses," an investor cares about the level of his consumption relative to the accustomed national average level of consumption (modeled as the level of national consumption per capita in the previous year). In this case, what an investor needs to guard against is not a decline in his own consumption per se, but a decline in consumption *relative* to the national level of consumption per capita attained in the previous year. With the level of consumption per capita generally growing over time, stocks that have a risk of occasional negative rates of return appear very risky; investors would be willing to hold stocks only if they offer a large expected equity premium. Using this modification of the utility function in simulations of the CCAPM can produce average rates of return of 6.70 percent per year on stocks and 2.07 percent per year on bills, with a value for A equal to only 6.[19]

In the other approach to modeling an accustomed level of consumption—known as "habit formation"—an individual investor's utility in any year depends on his level of consumption in that year compared to the level of his own consumption in the recent past.[20] Like the "Catching up with the Joneses" model, habit formation makes investors more loath to hold risky assets that could earn negative net rates of return. Thus, stocks will have to offer a sizable equity premium for investors to be willing to hold them in their portfolios. Abel (1990) and George Constantinides (1990) have used habit formation in the CCAPM with low values of A to generate fairly realistic values for the equity premium.

CONCLUSION

Rather than discouraging use of the CCAPM, the equity premium puzzle has provided the impetus for new lines of research aimed at making the statistical predictions of the CCAPM conform more closely to the statistical behavior of actual rates of return. One line of research has focused on producing additional data on asset returns and characterizing the statistical behavior of the actual rates of return on stocks and bills. This line of research has produced useful new information about the statistical properties of asset returns over an extended period of time.

Another line of research has focused on modifications of the basic CCAPM. Some of the modifications, such as taking account of differences among investors and incorporating more general attitudes toward risk, seem to help account for part of the large historically observed value of the average equity premium. But accounting for the equity premium is only a first step in accounting for the statistical behavior of asset returns. A good model of asset returns should also account for other statistical properties, such as the variability or

[19]These calculations are reported in Abel (1990).

[20]Another modification of the attitude toward risk is studied by Nason (1988), who introduces a time-varying lower bound on consumption in the utility function. This formulation has some analytic similarities to "Catching up with the Joneses" and habit formation, though it differs from these formulations.

predictability of returns.[21] In addition, a model that relates asset returns to consumption should be tested to see whether it is consistent with data on consumption by individuals and by the economy as a whole.

If incorporating differences among investors or more general attitudes toward risk can explain the various statistical properties of asset returns—and if the results are consistent with data on consumption—then the theories of both long-run economic growth and real business cycles will need to take account of these modifications.

[21]Some of the research discussed in this article, notably Cecchetti, Lam, and Mark (1990), Kandel and Stambaugh (1990, 1991), and Constantinides (1990), has already begun to examine other statistical properties of returns, but more remains to be studied.

REFERENCES

Abel, Andrew B. "Asset Prices Under Heterogeneous Beliefs: Implications for the Equity Premium," Working Paper 9-89, Rodney L. White Center for Financial Research, University of Pennsylvania (February 1989).

Abel, Andrew B. "Asset Prices Under Habit Formation and Catching Up With the Joneses," *American Economic Review* 80 (May 1990), pp. 38-42.

Cecchetti, Stephen, Pok-sang Lam, and Nelson Mark. "The Equity Premium and the Risk-Free Rate: Matching the Moments," mimeo, Department of Economics, Ohio State University (June 1991).

Constantinides, George M. "Habit Formation: A Resolution of the Equity Premium Puzzle," *Journal of Political Economy* 98 (June 1990), pp. 519-43.

Kahn, James A. "Moral Hazard, Imperfect Risk-Sharing, and the Behavior of Asset Returns," *Journal of Monetary Economics* 26 (August 1988), pp. 27-44.

Kandel, Shmuel, and Robert F. Stambaugh. "Expectations and Volatility of Consumption and Asset Returns," *Review of Financial Studies* 3 (1990), pp. 207-32.

Kandel, Shmuel, and Robert F. Stambaugh. "Asset Returns and Intertemporal Preferences," *Journal of Monetary Economics* 27 (February 1991), pp. 39-71.

Kocherlakota, Narayana R. "Disentangling the Coefficient of Relative Risk Aversion From the Elasticity of Intertemporal Substitution: An Irrelevance Result," *Journal of Finance* 45 (March 1990), pp. 175-90.

Mankiw, N. Gregory. "The Equity Premium and the Concentration of Aggregate Shocks," *Journal of Financial Economics* 17 (1986), pp. 211-19.

Mankiw, N. Gregory, and Stephen P. Zeldes. "The Consumption of Stockholders and Nonstockholders," *Journal of Financial Economics* (forthcoming 1991).

REFERENCES (continued)

Mehra, Rajnish, and Edward C. Prescott. "The Equity Premium: A Puzzle," *Journal of Monetary Economics* 15 (March 1985), pp. 145-61.

Mehra, Rajnish, and Edward C. Prescott. "The Equity Risk Premium: A Solution?" *Journal of Monetary Economics* 22 (July 1988), pp. 133-36.

Nason, James N. "The Equity Premium and Time-Varying Risk Behavior," Finance and Economics Discussion Paper 11, Board of Governors of the Federal Reserve System (February 1988).

Reitz, Thomas A. "The Equity Risk Premium: A Solution," *Journal of Monetary Economics* 22 (July 1988), pp. 117-31.

Siegel, Jeremy J. "The Real Rate of Interest From 1800-1990: A Study of the U.S. and the U.K.," Working Paper 9-91, Rodney L. White Center for Financial Research, University of Pennsylvania (March 1991).

Weil, Philippe. "The Equity Premium Puzzle and the Riskfree Rate Puzzle," *Journal of Monetary Economics* 24 (1989), pp. 401-21.

Weil, Philippe. "Equilibrium Asset Prices With Undiversifiable Labor Income Risk," mimeo, Harvard University (July 1990).

Understanding National and Regional Housing Trends

*Leonard Mills**

As the recession unwinds, housing starts will rebound from their current low levels. But over time, cyclical influences on housing starts will be overshadowed by the demographic factors that largely determine the trend in housing starts. The slowdown in adult population growth in the 1990s is a key factor in forecasts of a lower future level of housing starts.

This slowdown in population growth, however, won't be uniform across all regions; indeed, in some areas, it is expected to be quite pronounced. And so, given the strong link between population and housing, it is reasonable to expect the decline in the number of housing starts to affect some regions more than others.

Undoubtedly, there will be cyclical swings in housing in the years ahead, and these swings will affect regions differently. But policymakers, builders, and others concerned about the housing outlook should keep an eye on the slower population growth and its effect in lowering the number of housing starts.

*Leonard Mills is Manager of Financial Economics at the Federal National Mortgage Association and a former Senior Economist at the Federal Reserve Bank of Philadelphia.

LINKING POPULATION AND HOUSING

Housing researchers usually analyze the link between the adult population and the number of housing units in terms of two ratios: 1) the *vacancy rate*, the number of vacant houses divided by the total number of houses; and 2) the *headship rate*, the number of households divided by the adult population. These two components are affected by different economic and sociological factors.

Vacancy rates are partly affected by business-cycle conditions. For example, when income growth slows during a recession, people are less willing or able to afford the higher mortgage payments associated with new homes. To the extent that the recession is unanticipated by builders, the inventory of new homes—which are vacant homes—rises. The same is true for apartments; more of them become vacant during recessions.

The vacancy rate is also subject to long-term structural changes, such as laws that affect the cost of carrying a vacant housing unit. For example, any tax law change that accelerates depreciation deductions will lower the after-tax cost of carrying a vacant unit; stretching the deductions out over more years will raise the cost.[1] Any shifts in the vacancy rate due to tax law changes—or other structural changes—generally last longer than those due to the business cycle because tax laws change relatively infrequently.[2]

The other factor affecting the link between population and housing is the headship rate. The headship rate measures adults' tendencies to form households. For example, two adults could choose to live together to form one household and reside in the same housing unit. For these two adults, the headship rate would be 0.5. Alternatively, they could choose to live in two separate housing units. In this case, two households would be formed and the headship rate for the two adults would be 1.

Like vacancy rates, headship rates are subject to both business-cycle changes and long-term changes. When incomes are low during a recession, adults are more likely to join together to form a single household because they may be unable to afford living alone.[3] Thus, the headship rate has a tendency to fall during recessions. Longer-term changes in the headship rate include such factors as a fall in marriage rates or an increase in divorce rates, both of which decrease the tendency for adults to get together and form households.

The link between population and housing can be summarized by combining these two components. Specifically, the number of housing units per adult (HPA) can be computed as follows:

$$\frac{\text{Headship Rate}}{(1 - \text{Vacancy Rate})} = \text{Housing Units Per Adult (HPA)}$$

For example, the average U.S. vacancy rate for housing units since 1973 has been 3.2 percent,

[1]In 1981, depreciation deductions were accelerated, but five years later the depreciable life of residential real estate was lengthened. For a discussion of these tax law changes, see Stephen A. Meyer, "Tax Cuts: Reality or Illusion?" this *Business Review* (July/August 1983), and Theodore Crone, "Housing Costs After Tax Reform," this *Business Review* (March/April 1987).

[2]For a discussion of some other factors affecting the natural vacancy rate in rental housing units, see Stuart A. Gabriel and Frank E. Nothaft, "Rental Housing Markets and the Natural Vacancy Rate," *Journal of the American Real Estate and Urban Economics Association* 16 (1988), pp. 419-29.

[3]For a discussion of the effect of income and other factors on household formation decisions, see Lawrence B. Smith et al., "The Demand for Housing, Household Headship Rates, and Household Formation: An International Analysis," *Urban Studies* 21 (1984), pp. 407-14. Also see Patric H. Hendershott and Mark Smith, "Household Formations," in *The Level and Composition of Household Saving*, Patric H. Hendershott, ed. (1985).

and the headship rate—for the population 21 and over—has averaged 53.8 percent.[4] (See *Components of Housing Unit Per Adult*.) These two components lead to the calculation of 0.556 housing units per adult [0.538 ÷ (1-.032)]. In other words, there have been about 1.8 (1 ÷ 0.556) adults living in each housing unit.

Given that the number of housing units per adult has averaged 0.556 and that the 21-and-over population has increased by 39 million since 1973, the trend increase in the housing stock has been 22 million units (0.556 x 39 million). Meanwhile, the actual increase in the housing stock was 23 million units, a level close to the trend.

In summary, then, HPA is simply the link that allows one to translate population growth into housing growth.

The Link Between Population and Housing Has Been Stable. The arithmetic shown above is simple enough, but there is a complication: the link between population and housing may not be stable. The HPA will change whenever one of its two components changes.

[4]In this article, total housing units are defined as the sum of owner-occupied units, renter-occupied units, units for sale, and units for rent. Thus, units rented or sold but not occupied, units held for occasional use, and other vacancies are excluded from this analysis.

Components of Housing Unit Per Adult

For Total Housing Stock

	Nation	Northeast	Midwest	South	West
Averages (1973-87)					
Vacancy Rate	.032	.025	.030	.038	.033
Headship Rate	.538	.521	.541	.539	.554
Housing Units Per Adult	.556	.534	.558	.559	.573
Ranges for Housing Units Per Adult					
High (year)	.562 (87)	.542 (85)	.570 (81)	.576 (87)	.589 (78)
Low (year)	.550 (80)	.525 (73)	.548 (73)	.542 (83)	.560 (83)

For Owner-Occupied Housing Stock

	Nation	Northeast	Midwest	South	West
Averages (1973-87)					
Vacancy Rate	.015	.011	.014	.017	.016
Headship Rate	.602	.580	.609	.604	.619
Homeowner Rate	.646	.603	.689	.669	.600
Housing Units Per Adult	.395	.354	.426	.411	.377
Ranges for Housing Units Per Adult					
High (year)	.405 (81)	.366 (81)	.444 (81)	.421 (79)	.393 (76)
Low (year)	.388 (87)	.339 (73)	.408 (87)	.402 (83)	.357 (87)

An increase in the vacancy rate means that there are relatively more houses with no one living in them, which raises the number of houses per adult. HPA will also vary with the headship rate. A higher headship rate is associated with a higher number of housing units per adult, and a lower headship rate is associated with a lower number of units per adult. Because both the vacancy and headship rates depend on a variety of economic and sociological factors that can change over time, the HPA can also change over time. Accordingly, a constant HPA, such as the historical average used above, might not provide an accurate assessment of population-related trends in housing.

In fact, the HPA has varied only slightly since 1973, ranging from a high of 0.562 in 1980 to a low of 0.550 in 1987. Thus, the high and the low for HPA are only about 1 percent from the historical average of 0.556. More important, there has been no noticeable trend in the HPA since 1973.

Since the link between housing and population seems fairly stable, the trend in housing can be reliably computed by multiplying the average HPA by the 21-and-over population. This population-driven trend in housing is interpreted as the number of houses required by the size of the adult population (Figure 1). Of course, the actual housing stock has differed from this long-term trend. This difference occurs whenever HPA differs from its long-term average. For example, because the headship rate tends to fall during a recession, HPA tends to fall below its historical average and consequently the housing stock falls below its trend line.[5] In other words, any deviations in the actual housing stock from its population-driven trend line appear as temporary cyclical deviations.[6]

FIGURE 1
U.S. Housing Stock
(1973-2000)

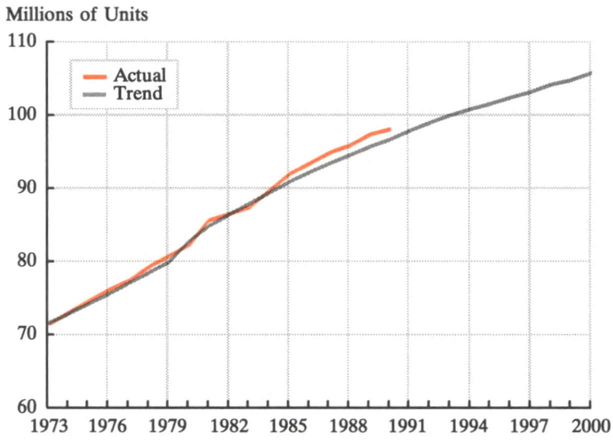

[5]The tendency for the residential vacancy rate to rise during a recession will raise the HPA, which offsets the effect of the declining headship rate during the recession. However, even though vacancy rates are subject to more cyclical variability than headship rates, the effect of the vacancy rate variability on HPA is generally smaller because the vacancy rate is so low.

[6]For statistical evidence on the use of the historical average HPA in projecting housing trends, see Theodore Crone and Leonard Mills, "Forecasting Trends in the Housing Stock Using Age-Specific Demographic Projections," *Journal of Housing Research* 2 (1991), pp. 1-20. This study found that, over the 1965-89 period, the HPA-based trend was more precise for owner-occupied units than for total units.

PROJECTING THE FUTURE HOUSING STOCK

By making some assumptions about the future size of the adult population and HPA, we can project the future housing stock. Since the bulk of the people who will make up the adult population have already been born, adult population projections for the next 10 to 20 years are considered very reliable. The Census Bureau regularly projects the future adult population by making assumptions about other determinants of the adult population, such as immigration and deaths.[7]

These other determinants are relatively easy to project over a span as short as a decade. For HPA, a useful baseline is to assume that the link between population and housing will remain as it has been. On the basis of historical evidence, using the average HPA has provided a reliable housing trend. For the 1990s, then, the trend growth in the housing stock is projected to slow because the adult population is projected to grow more slowly. The lower trend growth tells us that fewer houses will need to be added to the housing stock.

Less Need for Additional Housing Stock Means Lower Housing Starts. Changes in the housing stock occur when new housing units are built or old units are removed. Net additions to the housing stock are defined as the number of new housing starts minus the number of removals. In other words, in any given year, the number of housing starts must equal net additions to the stock plus removals.[8]

For the 1990s, the net additions required to accommodate the number of new households should average about 900,000 units per year. Adding the historical average of 300,000 removals per year results in an annual average of about 1.2 million housing starts in the 1990s. This is substantially below the annual level of housing starts in the 1970s and 1980s (Figure 2). Trend, as opposed to actual, housing starts are expected to fall from about 1.4 million in 1991

FIGURE 2
Housing Starts
Annual Averages

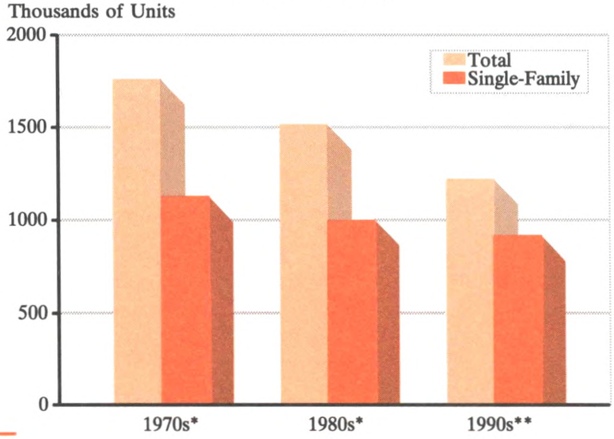

* Actual
** Trend

[7]Both the regional and national population projections used in this study are provided by the Census Bureau. The national projections are consistent with the "middle" (intermediate) series of the Census Bureau.

[8]Housing starts are measured directly, but removals must be measured indirectly by subtracting the number of housing starts in a given year from the actual change in the housing stock. Since 1973, removals have averaged about 300,000 units per year. Removals were higher in the 1960s because of the large number of urban renewal programs.

to about 1.1 million by the mid-decade.[9]

Single-Family Starts Are an Important Component. About two-thirds of the nation's households own their homes. Moreover, the owner-occupied segment of the housing stock is generally regarded as more stable because it seems less susceptible to waves of speculative investment. In addition, tax laws for owner-occupied housing have changed less frequently than those for rental housing. For these and other reasons, the owner-occupied segment of the housing market is often analyzed separately.

Because there is less turnover in the owner-occupied housing market, the vacancy rate is naturally lower than that for renter-occupied units. The owner-occupied vacancy rate has averaged 1.5 percent since 1973. Moreover, a very large proportion of home owners are in the 25-and-over age category, so that is the most relevant population segment for calculating the headship rate for owner-occupied units. Since 1973, the headship rate for the population 25 and over has averaged 60.2 percent.

In addition to headship and vacancy rates, the home-ownership rate is a necessary factor in analyzing the owner-occupied segment. The home-ownership rate is the number of home-owning households divided by the total number of households. The historical average for the home-ownership rate has been 64.6 percent. Multiplying the headship rate by the home-ownership rate, and dividing by 1 minus the vacancy rate, determines the number of owner-occupied houses per adult. Since 1973, the average has been 0.395, varying little.

A projection for single-family housing starts can be obtained by following the procedure outlined above for total housing starts.[10] First, multiply the change in the projected 25-and-over population by the average number of owner-occupied houses per adult, 0.395. Then add the historical average number of removals for the owner-occupied housing units (about 200,000).

On the basis of this calculation, single-family starts are also projected to decline in the 1990s. The trend in single-family starts falls to an annual average of 875,200 in the 1990s, compared to 1 million in the 1980s. The trend level of single-family starts falls from about 1 million in 1991 to slightly more than 740,000 by the end of the decade.

The population cohort born during the "birth dearth" period between 1972 and 1978 will, of course, enter the 21-and-over population before it enters the 25-and-over population. Consequently, the drop in the trend for single-family starts occurs later in the 1990s than the drop in the trend for total starts. In fact, the decline in the population-driven trend for single-family starts is not expected to be noticeable until 1996, while the decline in total starts will begin in 1993.[11]

Other Housing Market Aspects Are More Difficult to Predict. Even though the number of housing starts is expected to fall in the 1990s, that does not imply that the residential construction sector of the U.S. economy will decline. Although new housing measured in *physical units* (housing starts) will decline, each housing unit may be larger or more elaborate. And larger, more elaborate housing units are generally more expensive to build. Thus, new

[9]With the recession, actual housing starts through the first half of 1991 averaged only 956,000 at an annual rate. That the trend level of starts in 1991 is so much higher is one reason why many forecasters expect a rebound in housing activity once the recession ends.

[10]Single-family starts and additions to owner-occupied units are closely related because most occupied single-family units are owner occupied (85 percent in 1985) and few single-family units are built as rentals.

[11]For a broader discussion of the economic impact of the birth-dearth generation, see Theodore Crone, "The Aging of America: Impacts on the Marketplace and Workplace," this *Business Review* (May/June 1990).

housing measured in *dollar units* (residential investment) may not decline as much as housing starts. So residential construction workers and suppliers of residential building materials may not experience as large a drop in their business activity.

In addition, the trend in the size of the adult population tells us nothing about the long-term price of each housing unit.[12] Of course, a temporary overabundance of housing units relative to the population-driven trend, either regionally or nationally, will depress housing prices temporarily. Similarly, if the housing stock falls below the population-driven trend, there will be a tendency for house prices to be bid up. But these effects are only temporary because eventually the housing stock does return to its long-term trend. The long-term price of a housing unit appears to be determined not so much by changes in housing demand due to population trends, but by such factors as the cost of land, materials, and labor.[13]

HOUSING PROJECTIONS DIFFER BY REGION

The average HPA varies across different regions of the country. For the total housing stock, the Northeast has the lowest average HPA and the West has the highest. An implication of this regional difference in HPA is that even if the Northeast and West experience identical increases in their adult populations, more houses would have to be built in the West.

The regional differences in owner-occupied houses per adult are larger than those for the total housing stock because home-ownership rates also differ across regions. The Northeast also has the lowest average number of owner-occupied houses per adult. But the Midwest has the highest average number of owner-occupied houses per adult, primarily because it has the highest home-ownership rate.

Even if we consider only the adult segment of the population, regional populations are more difficult to project than the national population because we must also make assumptions about internal migration within the country. Using one set of Census Bureau assumptions regarding internal migration and assuming the number of houses per adult in each region equals its historical average, we can use regional population projections to predict the trend in housing starts by region.[14] Under this set of assumptions, the Midwest will experience the largest percentage decline in total housing starts, followed by the Northeast. This regional pattern holds for single-family starts as well, with the drop occurring later in all regions. (See *Projected Trends in Housing Starts*, p. 22.)

---

[12]This is the consensus among housing analysts. However, for an argument that population trends have a large impact on housing prices, see N. Gregory Mankiw and David N. Weil, "The Baby Boom, the Baby Bust, and the Housing Market," *Regional Science and Urban Economics* 19 (1989), pp. 235-58. Arguments against any impact of population trends on housing prices are more prevalent. For example, see Denise DiPasquale and William Wheaton, "Housing Market Dynamics and the Future of Housing Prices," Joint Center for Housing Studies of Harvard University, Working Paper W90-3; A. Steven Holland, "The Baby Boom and the Housing Market: Another Look at the Evidence," *Regional Science and Urban Economics* (forthcoming); and "Will Home Prices Collapse?" various special reports compiled by the National Association of Home Builders (1990).

[13]See James R. Follian, "The Price Elasticity of the Long-Run Supply of New Housing," *Land Economics* (1979), pp. 190-99.

[14]The Census Bureau's "Series A" for regional population projections is used in the housing projections. For a description of the assumptions underlying these projections, and alternative projections, see Signe I. Wetrogan, "Projections of Population of States by Age, Sex, and Race, 1989 to 2010," *Current Population Reports: Population Estimates and Projections*, Series P-25, No. 1053 (1990). Of course, regional housing projections will differ under different migration assumptions.

Projected Trends in Housing Starts
(Thousands)

Total Housing Starts

	Nation	Northeast	Midwest	South	West
1991	1415	158	167	621	470
1992	1438	164	168	642	465
1993	1259	127	113	594	425
1994	1138	105	78	554	401
1995	1072	95	61	528	387
1996	1107	104	70	537	395
1997	1062	99	58	519	385
1998	1126	111	73	543	398
1999	1131	115	75	545	395
2000	1186	126	88	565	407

Single-Family Housing Starts

	Nation	Northeast	Midwest	South	West
1991	1000	123	145	449	283
1992	950	115	126	433	276
1993	914	108	111	423	272
1994	928	107	111	432	277
1995	942	109	112	441	279
1996	968	112	118	457	280
1997	838	87	80	419	252
1998	753	72	57	389	236
1999	716	67	46	373	229
2000	743	72	54	382	234

Note: The sum of the regions may not add up to the nation because of rounding.

COULD THESE HOUSING PROJECTIONS BE WRONG?

By simply analyzing trends in the size of the adult population, we can construct a historically reliable baseline for the number of houses to be built in the 1990s. Inevitably, the projections will differ from the number of houses actually built, for a variety of reasons. Housing starts have a large cyclical component, which leads to deviations from trend. For example, housing starts so far in 1991 have been substantially below their trend level of 1.4 million units, partly because of the recession. And adjustments to these cyclical deviations may take some time to correct themselves. But over time, history leads us to believe that housing will return to its population-driven trend.

However, it *is* possible that the population-driven trend itself may change. First, national population projections must make some assumptions about immigration and other factors affecting the adult population. And regional projections must make additional assumptions about internal migration within the country. The assumptions underlying the population projections may turn out to be inaccurate, which in turn would make the housing projections inaccurate.

Second, long-term changes in headship, vacancy, and home-ownership rates may also occur, and such changes would alter the link between population and housing. For example, some analysts think the headship rate may rise slightly in the 1990s as the baby-boom generation ages further.[15] They argue that, historically, headship rates have continued to rise as people grow older, even though the increase after age 25 is very gradual. Since a higher headship rate would raise the HPA, the trend in housing starts may be higher in the 1990s. This argument assumes that the baby-boom generation will have the same tendencies to form households as their parents did. However, an alternative assumption is that each generation, including the baby-boomers, behaves uniquely and therefore long-term changes in headship rates are difficult to predict.

In any case, prospective changes in the vacancy, headship, and home-ownership rates will have very little impact on the trend in housing starts compared to the effect of the slower population growth. For example, even if headship rates were to rise by 10 percent—a figure far greater than any analyst expects—the annual average for total housing starts would rise only by 90,000, to about 1.3 million, still far below the level of starts in the 1970s and 1980s.[16] Thus, housing analysts, though they may disagree on the precise numbers, are unanimous in their belief that housing starts will decline substantially in the 1990s. Consequently, builders will build fewer houses in this decade, and policymakers and others concerned with the housing outlook should not be surprised by the lower number of starts.

[15]See, for example, Gretchen A. Armijo et al., "Demographic and Economic Trends," *Journal of Housing Research* 1 (1990), pp. 21-42.

[16]This calculation continues to assume that the vacancy rate and level of removals remain unchanged.

FEDERAL
RESERVE BANK OF
PHILADELPHIA

BUSINESS REVIEW Ten Independence Mall, Philadelphia, PA 19106-1574

Address Correction Requested