The Efficiency of the Market for Single-Family Homes

By KARL E. CASE AND ROBERT J. SHILLER*

The market for single-family homes does not appear to be efficient. Year-to-year changes in prices tend to be followed by changes in the same direction in the subsequent year. Moreover, information about real interest rates does not appear to be incorporated in prices. There is thus a profitable trading rule for persons who are free to time the purchase of their homes. Still, overall, individual housing price changes are not very forecastable.

There is good reason to think that the market for single-family homes ought to be less efficient than are financial markets. The market is dominated by individuals trading in the homes they live in. Because of transactions costs, carrying costs, and tax considerations, professionals find it relatively difficult to take advantage of profit opportunities in this market. For these reasons, it is commonly casually asserted that the market for single-family homes is inefficient, and "bull markets" in housing (i.e., temporary upward inertia in housing prices) are frequently alleged. But it is hard to find scholarly work confirming whether this is so.

We have found surprisingly little in the literature on the testing of the efficiency of real estate markets. A computer search turned up only three recent papers, by George Gau (1984–85) and Peter Linneman (1986). Gau describes his two papers as the "first rigorous testing" of real estate market efficiency."[1] His data, however, were confined to commercial real estate and to the Vancouver area for the years 1971–80. He concluded that prices in the Vancouver market were well described as a random walk. Linneman, who asserts that "there are no empirical studies of the efficiency of the housing market,"[2] did a study using observations on individual owner's assessments of house value (rather than actual sales prices) in Philadelphia for two points of time: 1975 and 1978. He found that houses that were undervalued relative to a 1975 hedonic regression (i.e., that had negative residuals in a regression of price on housing characteristics) tended to increase in value subsequently, but that because of transactions costs only an insignificant number of units appear to present profitable arbitrage candidates. Robert Engle, David Lilien, and Mark Watson (1985) estimated a model of the resale housing market using data on retail house sales in San Diego for 1973–80. They concluded that much of the overall movement in housing prices in this period could be explained in terms of such factors as demographically driven changes in the cost of housing services, proposition 13, and the inflation-driven change in marginal tax rates. But they did not investigate directly whether the market was efficient.

This paper performs tests of the efficiency of this market using data from the Society of Real Estate Appraisers' tapes for the years 1970 to 1986 for Atlanta, Chicago, Dallas, and San Francisco/Oakland; see the Appendix. The tapes contain actual sale prices and other information about the homes. We

*Department of Economics, Wellesley College, Wellesley MA 02181, and Department of Economics, Cowles Foundation, Yale University, Box 2125 Yale Station, New Haven, CT 06520, respectively. We are grateful to Catherine Christensen and Maura Doyle for excellent research assistance. Arnold Kling and participants at seminars at Brown University, Harvard University, and Princeton University made helpful comments. The research was supported by the Federal Reserve Bank of Boston and the National Science Foundation under grant no. SES-8717080. Data used in this paper (files of multiple sales) will be made available on request.

[1] George Gau, 1984, p. 301.

[2] Peter Linneman, p. 140.

extracted from the tapes for each city a file of data on houses sold twice for which there was no apparent quality change and for which conventional mortgages applied. For each house the data we used consisted of the two sales prices and the two dates (quarters) in which the sales occurred. The total number of observations on such double sales of relatively unchanged homes was 39,210 (8,945 Atlanta, 15,530 Chicago, 6,669 Dallas, and 8,066 San Francisco). None of the other studies had actual repeat sales price data on individual homes at all, let alone such a large number, and none of the studies spanned the time interval and geographical area of our study. Moreover, the present study presents some statistical-methodological improvements over the Gau studies in its effort to test the random-walk theory for housing prices.

I. The WRS Index

In a companion paper we discuss our method of price index construction, which we call the Weighted Repeated Sales (WRS) method. The method is a modification of the regression method proposed by Martin J. Bailey, Richard Muth, and Hugh Nourse (hereafter, BMN). The BMN method produces estimates and standard errors for an index of housing prices by regressing, using ordinary least squares, the change in log price of each house on a set of dummy variables, one dummy for each time period in the sample except for the first. Each value of the log price index BMN(t) is represented by a regression coefficient, except for the first value of the log price index, which is set to zero as a normalization. The dummy variables are zero except that the dummy is $+1$ corresponding to the second time period when the house was sold and that the dummy is -1 corresponding to the first time period when the house was sold (unless this is the first time period). Bailey et al. argued that if the log price changes of individual houses differ from the citywide log price change by an independent, identically distributed noise term, then by the Gauss-Markov theorem their estimated index is the best linear unbiased estimate of the citywide log price.

Our procedure differs from the BMN procedure because we feel that the house-specific component of the change in log price is probably not homoscedastic but that the variance of this noise increases with the interval between sales. The motivation for our WRS method was the assumption that the log price P_{it} of the ith house at time t is given by

$$P_{it} = C_t + H_{it} + N_{it}, \tag{1}$$

where C_t is the log of the citywide level of housing prices at time t, H_{it} is a Gaussian random walk (where ΔH_{it} has zero mean and variance σ_h^2) that is uncorrelated with C_T and H_{jT} $t \neq j$ for all T, and N_{it} is an identically distributed normal noise term (which has zero mean and variance σ_N^2) and is uncorrelated with C_T and H_{jT} for all j and T and with N_{jT} unless $i = j$ and $t = T$. Here, H_{it} represents the drift in individual housing value through time, and N_{it} represents the noise in price due to imperfections in the market for housing.[3] Presumably, the value that a house brings when it is sold depends on such things as the random arrival of interested purchasers, the behavior of the real estate agent, and other random factors, so that the sale price is not identical to true value. Moreover, there may be some change in true value that may be bunched at the purchase date.

A three-step weighted (generalized) least squares procedure was undertaken. In the first step, the BMN procedure was followed exactly, and a vector of regression residuals was calculated. In the second step, the squared residuals in the first step regression were regressed on a constant and the time interval between sales.[4] The constant term was the estimate of σ_N^2, and the slope term was the estimate of σ_H^2. In the third step, a generalized least squares regression (a weighted regression) was run by first divid-

[3] Cary Webb has a similar model, except that $\sigma_N = 0$.

[4] Because the errors in this regression are likely to be larger for houses for which time interval between sales is larger, a weighted regression was used, downweighting the observations corresponding to large time intervals.

ing each observation in the step-one regression by the square root of the fitted value in the second-stage regression and running the regression again.

The estimated WRS index WRS(t) and its accuracy are discussed in our companion paper. The level of the index is quite well measured, the quarterly first difference of the index is not well measured, and the annual difference of the index is fairly well measured. One way of describing how well these variables are measured is to compute the ratio of the standard deviation of a variable to the average standard error for that variable. For the log index in levels, this ratio is 13.87 for Atlanta, 24.52 for Chicago, 9.94 for Dallas, and 28.03 for San Francisco-Oakland. Thus, we can make very accurate statements about the levels of house prices in the cities. For the quarterly difference of the log indexes, the ratio is 1.64, 1.61, 1.35, and 1.54, respectively. We thus cannot accurately describe the quarterly changes in the log prices, though the index will give a rough indication. For the annual difference of the log index, the ratio is 2.73, 3.99, 2.90, and 3.62, respectively; we can make fairly accurate statements about the annual change in log housing prices.

Other existing housing price indexes are widely interpreted as showing even monthly changes in housing prices. We argue in our companion paper that these indexes (for which no standard errors are provided) are likely to be less accurate than ours.

II. Seasonality, City Influence, and Beta

Part A of Table 1 gives sample statistics for $W(t) = \mathrm{WRS}(t) - \ln(\mathrm{CPI}(t))$. $W(t)$ is the real WRS index in each city, deflated by the city-specific consumer price index. The growth in real price was less than 1 percent per quarter for all cities, even San Francisco where a real estate "boom" took place. The standard deviation in quarterly real price changes is less than 3 percent per quarter, or on the order of a third of the standard deviation of quarterly changes in comprehensive real stock price indexes.

Individual housing prices are like many individual corporate stock prices in the large standard deviation of annual percentage change, close to 15 percent a year for individual housing prices. But housing prices in our sample differ from stock prices in that the individual prices are not so heavily influenced by the aggregate market price. When 1-year changes in real individual house prices are regressed on contemporaneous 1-year changes for the real WRS index, the R squared is only 0.066 for Atlanta, 0.158 for Chicago, 0.121 for Dallas, and 0.270 for San Francisco.

While second quarter price changes tend to be high and third quarter changes low, the difference is small, and only in Chicago is seasonality statistically significant at the 5 percent level. The National Association of Realtors' series on the median price of existing single-family homes appears to show more pronounced seasonality; we argued elsewhere that much of this may be due to seasonality in the composition of houses sold over the year (see our companion paper). Still, the NAR and WAR indexes do agree that prices are highest midyear (the NAR index tends to peak in July).

The beta (estimated for each of the cities by regressing the quarterly change in the log nominal WRS index on the corresponding change in the log Standard & Poor's Composite Index) is always virtually zero (Table 1, Part B). This confirms the results of Gau.

III. Testing for Market Efficiency

One might think that we could test the random-walk property of prices by regressing the change in the index on lagged changes in the index. But there is a problem, the noise in the estimated index. To see this point, consider the very simple case where we have two observations only on log housing prices. House A was sold in period 0 and period 1, while house B was sold in period 0 and period 2. The estimated changes in the log price index (using either the original BMN or WRS procedure, since in this example the number of observations equals the number of coefficients) are, for period 1, $P_{A1} - P_{A0} = C_1 - C_0 + H_{A1} - H_{A0} + N_{A1} - N_{A0}$, and for period 2 $-(P_{A1} - P_{A0}) + (P_{B2} - P_{B0}) = C_2 - C_1 - (H_{A1} - H_{A0} + N_{A1} - N_{A0}) + H_{B2}$

TABLE 1—SUMMARY STATISTICS

A. Quarterly Changes in Real WRS Log Price Index: $z = W(t) - W(t-1)$

	All Quarters	Mean z for Quarter t				H_0: All Quarters Same Mean
	Mean z std. z	$t=1$ (t-stat)	$t=2$ (t-stat)	$t=3$ (t-stat)	$t=4$ (t-stat)	F Prob.
Atlanta	0.0001	−0.0013	0.0050	−0.0043	0.0006	0.33
	0.0270	(−0.2040)	(0.7694)	(−0.6461)	(0.0888)	0.85
Chicago	0.0007	0.0088	0.0071	−0.0019	−0.0115	3.32
	0.0218	(1.6456)	(1.3682)	(−0.3571)	(−2.1970)	0.02
Dallas	0.0050	0.0031	0.0114	0.0024	0.0028	0.43
	0.0265	(0.4612)	(1.7788)	(0.3586)	(0.4172)	0.79
San Fran.	0.0092	0.0100	0.0161	0.0024	0.0082	0.84
	0.0254	(1.5040)	(2.5822)	(0.3621)	(1.2317)	0.51

B. Regression of Nominal WRS Index Changes on Changes in Log Standard and Poor Composite Index:

$$WRS(t) - WRS(t-1) = \alpha + \beta(LSP(t) - LSP(t-1)) + u(t)$$

City	No obs. S.E.E.	α (t)	β (t)	R^2 $\bar{R}^2$
Atlanta	65	0.017	−0.022	0.003
	0.025	(5.264)	(−0.454)	−0.013
Chicago	65	0.017	−0.014	0.002
	0.018	(7.418)	(−0.393)	−0.013
Dallas	65	0.023	−0.066	0.026
	0.027	(6.698)	(−1.289)	0.010
San Fran.	66	0.025	0.035	0.006
	0.028	(7.259)	(0.643)	−0.009

Note: WRS(t) is the quarterly WRS index (in logs) described in the text, $W(t)$ is WRS(t) deflated by the city-specific consumer price index averaged over the quarter. LSP(t) is the log of the Standard & Poor's Composite Index, quarterly average of daily prices. Sample is 1970-second quarter to 1986-second quarter (65 observations), except for San Francisco where the data are 1970-second quarter to 1986 third quarter (66 observations).

$-H_{B0} + N_{B2} - N_{B0}$. The index change between 0 and 1 is negatively correlated with the change between 1 and 2 because of common terms appearing with opposite signs.

There may also be positive serial correlation of estimated changes in the log price index. Suppose we have three houses in our sample, house A was sold in periods 1 and 3, house B was sold in periods 0 and 2, and house C was sold in periods 0 and 3. The estimated changes in the log price index (again, using either the original BMN procedure or the WRS procedure with, the full sample) are, for period 1, $(P_{C3} - P_{C0}) - (P_{A3} - P_{A1})$ and for period 3, $(P_{C3} - P_{C0}) - (P_{B2} - P_{B0})$. These two estimated changes will be positively correlated in our model because house C appears with the same sign in both expressions, while the specific shocks to the other two houses are independent. The three-house example also makes clear that there may be serial correlation between noncontiguous price changes.

Gau's procedure for testing the efficiency of the Vancouver commercial real estate market involved building three price indices (not repeat-sales indexes): sales price divided by square footage, sales price divided by gross income, and sales price divided by number of suites. For each month he chose a single transaction for his index. His method of construction of a price series is likely to induce the same spurious serial correlation in price changes. His conclusion that his

price index was approximately a random walk might be spurious.[5]

IV. A Simple Expedient for Dealing with Estimation Error

We have seen that we cannot test efficiency of the housing market by regressing real changes in the WRS index onto lagged changes, and testing for significance of the coefficients, because the same noise in individual house sales contaminates both dependent and independent variables. A simple expedient for dealing with this problem is to split the sample of individual house sales data and estimate two WRS indexes. For each city, houses were randomly allocated between samples A and B, and log price indexes WRS_A and WRS_B were estimated using the respective samples. Then efficiency is tested by regressing changes in the real log index $W_A(t) = WRS_A(t) - \ln(CPI(t))$ on lagged changes in the real index $W_B(t) = WRS_B(t) - \ln(CPI(t))$, where $CPI(t)$ is the consumer price index for the city for quarter t (quarterly average).[6] Both sides of the equation are contaminated by noise, but since the same houses do not enter into the indexes on the two sides of the equation, these noise terms will not be correlated. If the slope coefficients are statistically significant, we can reject weak form efficiency.

Table 2 presents such regressions. For each city, we report first the regression of annual change with real log index A on the contemporaneous annual change in real index B, as a diagnostic on our methods. The coefficient should be 1.00 if the indexes were measured perfectly, but should tend to be less than one for estimated indexes, due to the errors-in-variables problem. Fortunately, the estimated coefficients are never too far below 1.00. For each city, we then report the regression of the real annual change in the real index for sample A on the 1-year-lagged real annual change in the real index for sample B, and then the same regression with samples A and B reversed. These coefficients are always positive and substantial, and statistically significant at the 5 percent level for Chicago and San Francisco. The greater significance in Chicago may be due to the greater number of observations on individual houses for that city, so that the measurement error problem is less severe.

We interpret these results as substantial evidence that there is inertia in housing prices, increases in prices over any year tending to be followed by increases in the subsequent year.

The Table 2 regressions show that real price changes are forecastable, but do not show that there are any predictable excess returns to be had in investing in real estate. It is in principle possible that the forecastability of price changes is due to nothing more than the forecastability of real interest rates or of the dividend on housing. Table 3 reports analogous regressions, where the dependent variable is the after-tax excess nominal return on housing over the 1-year treasury bill rate, using one index, and the independent variable is the after-tax excess nominal return using the other index. The after-tax rate for sample A or B was defined by:

$$\text{Excess}_j(t) = \frac{\exp(WRS_j(t+4) - WRS_j(t)) + C_j[R(t)+R(t+1)+R(t+2)+R(t+3)]}{\exp(WRS_j(t))} - 1 - (1-\tau)r(t)/100 \qquad j = A, B$$

where $WRS_j(t)$ is the nominal (uncorrected for inflation) WRS index (in logs) estimated using sample j, $R(t)$ is the city-specific index, residential rent, from the U.S. Bureau of Labor Statistics, τ is the marginal

[5] It should be noted that a strength of Gau's approach relative to ours is that he could research the properties more thoroughly. He used detailed description of debt liens from provincial land title records, to adjust for financing with below-market interest rates. We did not have such information on the Society of Real Estate Appraisers' (SREA) tapes. He also controlled for other quality differences by his choice of properties to include.

[6] Since quarterly data were used and price index changes were measured over four quarters, error terms in the regression are not independent under the random-walk assumption, but follow an MA-3 process. A method of Lars Hansen and Robert Hodrick was used to correct the standard errors of the ordinary least squares estimates.

TABLE 2—REGRESSION OF CHANGES IN REAL LOG INDEX ESTIMATED WITH ONE-HALF OF SAMPLE ON CHANGES IN REAL LOG INDEX ESTIMATED WITH OTHER HALF OF SAMPLE

$$W_j(t) - W_j(t-4) = \beta_0 + \beta_1(W_k(t-L) - W_k(t-4-L)) + u(t)$$

t = 1972-I to 1986-II (1986-III San Francisco)

City Parameters	No. obs. S.E.E.	β_0 (t)	β_1 (t)	R^2 $\bar{R}^2$
Atlanta				
$j = A, k = B, L = 0$	58	0.001	0.857	0.629
	0.028	(0.074)	(5.981)	0.622
$j = A, k = B, L = 4$	58	−0.003	0.215	0.045
	0.045	(−0.279)	(0.991)	0.028
$j = B, k = A, L = 4$	58	−0.004	0.191	0.046
	0.041	(−0.408)	(1.051)	0.029
Chicago				
$j = A, k = B, L = 0$	58	−0.001	0.871	0.816
	0.024	(−0.208)	(9.337)	0.833
$j = A, k = B, L = 4$	58	−0.001	0.412	0.183
	0.053	(−0.076)	(1.953)	0.169
$j = B, k = A, L = 4$	58	−0.000	0.502	0.234
	0.054	(−0.011)	(2.226)	0.220
Dallas				
$j = A, k = B, L = 0$	58	0.002	0.730	0.658
	0.029	(0.317)	(6.264)	0.652
$j = A, k = B, L = 4$	58	0.011	0.254	0.090
	0.047	(0.857)	(1.474)	0.074
$j = B, k = A, L = 4$	58	0.012	0.312	0.046
	0.052	(0.874)	(1.460)	0.029
San Francisco				
$j = A, k = B, L = 0$	59	0.017	0.608	0.313
	0.063	(0.947)	(3.061)	0.301
$j = A, k = B, L = 4$	59	0.030	0.255	0.055
	0.074	(1.435)	(1.093)	0.038
$j = B, k = A, L = 4$	59	0.021	0.430	0.220
	0.062	(1.206)	(2.462)	0.206

Note: Houses were randomly allocated into two separate samples of half-original size, samples A and B. $W_A(t)$ is the real WRS index estimated using sample A only, $W_B(t)$ is the real WRS index estimated using sample B only. Both series are deflated using the city-specific consumer price index.

personal income tax rate (assumed to be 0.30), and $r(t)$ is the 1-year treasury bill rate, secondary market.[7] The constant C_j was chosen to make the average "dividend-price ratio" $C_j[R(t)+R(t+1)+R(t+2)+R(t+3)]/\exp(WRS_j(t))$ equal to 0.05. We are using the residential rent index to indicate the implicit "dividend" (in the form of housing services) on houses, and must guess as the factor of proportionality between the index and the actual dividend. The assumptions about taxes are that neither the capital gain nor the implicit rent are subject to income taxes, but that interest is deducted from taxable income.[8]

[7] The residential rent index is computed by the U.S. Bureau of Labor Statistics every other month only. For quarters where two months are available, $R(t)$ is the average of the two figures. When only one month is available, $R(t)$ is the figure for the middle month. The interest rate $r(t)$ is the quarterly average of the monthly series Treasury bills, secondary market, 1-year, from the Board of Governors of the Federal Reserve System.

[8] We should properly also account for changes through time in the property tax rate. However, existing data series do not appear to allow us to measure well changes in this rate for the cities and sample period studied.

TABLE 3—REGRESSION OF AFTER-TAX EXCESS RETURNS ESTIMATED WITH ONE-HALF OF SAMPLE ON AFTER-TAX EXCESS RETURNS ESTIMATED WITH OTHER HALF OF SAMPLE

$$Excess_j(t) = \beta_0 + \beta_1 Excess_k(t-L) + u(t+4)$$

City Parameters	No. obs. S.E.E.	β_0 (t)	β_1 (t)	R^2 $\bar{R}^2$	Trading Profits
Atlanta					
$j=A, k=B, L=0$	58	0.012	0.831	0.673	
	0.030	(1.036)	(6.171)	0.667	
$j=A, k=B, L=4$	58	0.041	0.327	0.113	0.009
	0.049	(2.159)	(1.556)	0.097	
$j=B, k=A, L=4$	58	0.038	0.348	0.135	0.010
	0.041	(2.141)	(1.782)	0.120	
Chicago					
$j=A, k=B, L=0$	58	0.004	0.915	0.862	
	0.026	(0.452)	(9.848)	0.859	
$j=A, k=B, L=4$	58	0.020	0.661	0.449	0.021
	0.052	(1.086)	(3.577)	0.439	
$j=B, k=A, L=4$	58	0.017	0.706	0.479	0.024
	0.051	(0.959)	(3.774)	0.470	
Dallas					
$j=A, k=B, L=0$	58	0.010	0.856	0.762	
	0.036	(0.735)	(7.555)	0.757	
$j=A, k=B, L=4$	58	0.037	0.526	0.286	0.014
	0.061	(1.570)	(2.778)	0.273	
$j=B, k=A, L=4$	58	0.038	0.549	0.286	0.017
	0.063	(1.550)	(2.737)	0.273	
San Francisco					
$j=A, k=B, L=0$	59	0.029	0.759	0.461	
	0.082	(0.991)	(3.881)	0.451	
$j=A, k=B, L=4$	59	0.055	0.507	0.203	0.024
	0.100	(1.502)	(2.130)	0.189	
$j=B, k=A, L=4$	59	0.046	0.550	0.379	0.029
	0.079	(1.708)	(3.474)	0.368	

Notes: Houses were randomly allocated into samples A and B. $Excess_A(t)$ is the city excess return estimated using sample A only, $Excess_B(t)$ is the city excess return estimated using sample B only. Rental index (used to compute returns) was scaled so that average dividend-price ratio was 0.05. Assumed income tax rate was 0.30. T = 1971-I to 1985-II (1985-III San Francisco). Trading profits are average extra return for the sample (A or B) as a proportion of the value of the house for the trading rule in the text.

As seen in Table 3, excess returns are even more forecastable than real price changes. The greater forecastability holds up even when we adjust the constant C_j to make the average dividend-price ratio either 0.0 or 0.1, adjust the tax rate τ up to 0.50, and whether we substitute the residential mortgage rate for the interest rate $r(t)$. Apparently, the greater forecastability of excess returns comes about largely because of the forecastability of real interest rates over this period, and because housing prices do not take account of information about predicted real interest rates. That real interest rates are quite forecastable may surprise some readers, who remember Eugene Fama's assertion that real interest rates are almost unforecastable. Fama's sample period in that paper was 1953 to 1971, which hardly overlaps with our sample period. Since 1971 real interest rates have shown major persistent movements and have been much more forecastable. Real interest rates shifted from positive to negative in the early 1970s, and sharply shifted up to large positive values following the October 1979 changes in the operating procedures of the Federal Reserve System (see John Huizinga and Frederic Mishkin, 1986). The forecastability of real interest rates is likely to have more impact on the forecastability of excess returns in citywide housing returns over the riskfree

rate than on the excess returns between corporate stock indexes over the riskfree rate, just because the variability of corporate stock price indexes is so much higher than the variability of citywide housing price indexes.

V. Exploiting Serial Dependence in the Housing Market

Observed serial dependence, of course, does not by itself imply that a market is inefficient in the full sense of the term. There must be some way of exploiting that dependence. The institutional structure of the housing market makes it appear at first glance that exploitation would be difficult. First, there is no market for futures contracts and there are no short sales so that there are no profit opportunities to exploit if the market is expected to decline. Second, transactions costs are high. Selling real estate traditionally involves a brokerage commission, typically 6 percent, which covers sellers' search costs. Since the product is heterogeneous, buyers incur high search costs. For a portfolio investor to realize gains from appreciation, a capital gains tax must be paid. For those buying a property to live in, there are moving costs.

The absence of a futures market or short sales means only that forecast decreases in home prices cannot be exploited. If excess returns are expected to be positive because of appreciation, there is nothing to preclude a buy-and-hold strategy.

Those interested in exploiting potential positive excess returns fall into three categories: (1) first-time home buyers who intend to live in the unit; (2) home owners who live in their units who desire to trade up, or increase the size of their holding; and (3) those who buy and sell properties in which they do not intend to live as portfolio investments. The institutional impediments to exploiting accurately forecast excess returns are different in case 3 than they are in cases 1 and 2.

Buyers in case 1 pay no brokerage fees, since these are borne by sellers, and no capital gains tax as long as they remain home owners. A buyer who wants to execute a purchase can also do so very quickly. Purchase and sale agreements can be negotiated in a matter of minutes if buyer and seller agree on price. There is, of course, a lag between purchase and sale and closing, but since price is fixed at time of purchase and sale, appreciation during the closing period accrues to the buyer. In case 2, the buyer is also a seller and must pay a brokerage fee on the sale of her earlier residence. Again, however, no capital gains tax liability is normally incurred.

Transactions costs for the portfolio investor of the speculator are higher since costs are incurred at the time of purchase and at the time of sale. At point of sale, the portfolio investor must pay a capital gains tax—during the period of this study, the maximum effective tax rate on capital gains was 20 percent.

To explore the potential for actually exploiting profit opportunities revealed by the forecasting regression in Table 3, we simulate a simple trading rule. Consider a first-time home buyer (case 1) or a home owner who wants to trade up (case 2), and suppose that this buyer is indifferent, given personal and financial considerations other than forecastable variations in excess returns, between buying now or waiting another year. We establish the following trading rule: buy now if the excess return predicted by the regression in Table 3 is greater than the mean excess return, otherwise delay purchase for 1 year. Implementing this rule does not require estimates of the parameters in Table 3 beyond the sign of the slope coefficient. Only the mean of the independent variable, the excess return, is needed. In this case we can disregard transactions costs since the purchase will be made anyway. We can also disregard capital gains taxes since we are assuming the gains are on a principal residence and they will be continuously rolled over. The rightmost column of Table 3 gives the average trading profits (as a proportion of the value of the house) individuals would have earned had they followed this trading rule over the sample period. The average trading profit is the proportion of quarters in the sample that are early buys given the trading rule times the difference between the mean excess returns in those periods where

an early buy is signaled by the regression and the mean excess return for the entire sample. The average trading profit can also be viewed as approximately equal to one-half times the mean of the "rectified" excess return. For a given quarter, rectified excess return is equal to the actual excess return if the trading rule gave a buy-early signal, otherwise equal to minus excess return.

The average trading profits are between 2.4 and 2.9 percent for San Francisco, which suggests that it would have been worthwhile for some potential purchasers to attend to the trading rule. This is the average over all quarters, and of course for roughly half the quarters, when an early buy was indicated, the average excess profits were twice as high. (Sometimes the predicted excess return differed the mean excess return by over 10 percentage points.) For Chicago and Dallas, the average trading profits fell in the 1.4 to 2.4 percent range. For Atlanta, the city with the least volatile prices, the average trading profits were only around 1 percent. But even here, there were quarters where the forecasting regression would predict much higher excess profits. (Sometimes the predicted excess return deviated from the mean excess return by over 4 percentage points.)

VI. Forecasting Individual House Sales Data

A second procedure for testing the efficiency of the market for single-family homes is to regress changes in *individual* housing prices between time t and a subsequent period on information available at time $t-1$. The log price index we construct appears only as an explanatory variable in these regressions, and so any spurious serial correlation in it will have no effect on our results. Under the efficient markets hypothesis, anything in the information set at time t should have no explanatory power for individual house price changes subsequent to that date. It is natural to set up the testing of the efficient markets hypothesis in this way: we are concerned with forecasting individual housing prices and if people were to use past price data to forecast these prices, the forecasting variable would be an index like ours.

To assure that the individual price changes are predicted using lagged information, we reestimated the WRS index anew for each quarter, using only data available up to that quarter. That is, we reestimated the entire WRS index for all N quarters in each sample, thus providing N different estimated price indexes, with from 1 to N time periods. In our forecasting regressions where past price indexes were used as explanatory variables, only those past values in the price index were used that were estimated using data up to and including the quarter before the quarter of the first sale of the house.[9]

Doing regression tests of the efficient markets hypothesis by regressing individual house log price changes does have a potential problem in that many of the observations of price changes are for time intervals that overlap with each other. Thus, we cannot assume that residuals are uncorrelated with each other, even if they are uncorrelated with the independent variables.

To deal with this overlap problem, we use the model (1) again where the null hypothesis of market efficiency is taken to be that C_t is a random walk that is independent of anything in the information set at time $t-1$. Consider two different houses in a city, house A sold at time t and t' and house B sold at time T and T'. The variance of the residual in the regression of the log real price on lagged information (under the null hypothesis of market efficiency this residual is just the change in price) for house A is $(\sigma_C^2+\sigma_A^2)(t'-t)+2\sigma_M^2$, and the covariance between the residual for house A and for house B is $n\sigma_C^2$, where n is the length of overlap of the two time intervals. The testing procedure was as follows. A preliminary or-

[9] Note that all three steps of the WRS estimation procedure were run separately for each quarter, using only data available in that quarter, so that no future information would creep into the constructed price index. In some instances (especially for the earlier quarters, that is, using small amounts of data) the step 2 estimated coefficient of the interval between sales had the wrong sign. When this happened, it was set to zero, so that the procedure then reduces to ordinary least squares in step 3.

TABLE 4—INDIVIDUAL HOUSE LOG PRICE CHANGES ON LAGGED REAL INDEX CHANGE

$$P(i,t_i+4)-P(i,t_i)=\beta_0+\beta_1(W(t_i-1,t_i-1)-W(t_i-1,t_i-5))+u(t+4)$$

City	No. obs. S.E.E.	β_0 (t-stat)	β_1 (t-stat)	R^2
Atlanta	246	0.0380	0.2392	0.002
	0.141	(2.6875)	(0.6155)	
Chicago	596	0.0416	0.3437	0.012
	0.137	(2.261)	(1.0588)	
Dallas	202	0.0874	0.0761	0.001
	0.146	(3.7157)	(0.2268)	
San Francisco	332	0.1000	0.3337	0.028
	0.125	(3.183)	(1.0108)	

Notes: In the regressions, each observation i corresponds to a house that was sold twice 4 quarters apart, and t_i denotes the quarter of the first sale for house i. Prices are in real terms: $P(i,t)$ is the natural log price of the ith home at time t minus the natural log of the city consumer price index for time t. $W(t,t')$ $t'<t$ is the WRS log price index for time t' estimated with data up to time t and minus the natural log of the city consumer price index for time t'. Figures in parentheses are t-statistics computed taking into account the serial correlation of error terms induced by overlapping intervals between sales.

dinary least squares regression (where $t'-t$ was fixed at a constant for all observations in the regression) was performed to get a vector of estimated residuals. The parameter $(\sigma_C^2+\sigma_A^2)(t'-t)+2\sigma_M^2$ was estimated as the average square value of the residuals. The parameter σ_C^2 was estimated by forming all possible products of residuals for different houses where the time intervals overlap, dividing each by the length of the overlap, and forming the average of these. The variance matrix Ω was constructed using these estimates, and the variance matrix of the ordinary least squares estimates was taken as $(X'X)^{-1}X'\Omega X(X'X)^{-1}$. This variance matrix was used to construct t-tests and chi-squared tests of market efficiency.

VII. Results with Individual House Data

The regression results generally do not find statistical significance (Tables 4 and 5). The magnitudes of coefficients estimated in Table 4 are however roughly consistent with those found in Table 2, and the distributed lag pattern in Table 5 shows a crude indication of an exponential decay pattern that gives most weight to the most recent quarterly index change. There appears to be a substantial response in individual house prices to lagged index changes, but there is so much noise in individual houses (the standard deviation of annual price changes is comparable to that on the aggregate stock market) that we do not generally find statistical significance.

One reason that the regressions did not disclose stronger or more consistent evidence of inertia in housing prices is inadequate data. While we had hundreds of observations of individual house sales for each forecast horizon, we have only 16 years of data. The serial correlation correction in effect does not assume a great number of "degrees of freedom" despite the large number of observations.

Errors in the WRS index as a measure of citywide prices are a problem tending to bias our coefficients, probably toward zero. The index is reestimated anew every quarter, and there is always substantial measurement error in the most recent observations of the index.[10]

[10] For example, with the San Francisco-Oakland data, there is, when the index is estimated with data through 1980-2, an estimated decline in real housing prices of 6.20 percent between 1980-1 and 1980-2 (the actual decline, not an annualized rate). When data through 1986-3 are used to estimate, the index between those two quarters is estimated to increase 2.53 percent.

TABLE 5—REGRESSIONS OF REAL LOG PRICE CHANGE ON LAGGED INDEX CHANGES

$$P(i,t_i+4)-P(i,t_i)=\beta_0+\Sigma(j=1,\ldots,4)\beta_j(W(t_i-1,t_i-j)-W(t_i-1,t_i-j-1))+u(i,t+4)$$

City	N.	β_0	β_1	β_2	β_3	β_4	S.E.E.	R^2	χ^2
Atlanta	246	0.037	0.432	0.283	−0.009	−0.029	0.142	0.008	1.154
		(2.919)	(1.033)	(0.602)	(−0.019)	(−0.075)			
Chicago	596	0.044	1.055	0.663	−0.253	−0.149	0.136	0.032	7.692
		(2.494)	(2.254)	(1.309)	(−0.565)	(−0.296)			
Dallas	202	0.089	0.430	0.220	0.094	−0.483	0.145	0.019	3.259
		(4.841)	(0.992)	(0.487)	(0.213)	(−1.172)			
SF/Oak.	332	0.099	0.652	0.511	0.118	−0.106	0.125	0.036	2.822
		(3.325)	(1.465)	(1.173)	(0.222)	(−0.214)			

Notes: χ^2 is chi-squared statistic (4 degrees of freedom) for null hypothesis that all slope coefficients are zero. The chi-square tests are computed taking into account the serial correlation of error terms induced by overlapping intervals between sales. See also notes to Table 4 above.

To attempt to deal with this problem, a time-varying errors-in-variable model was estimated. It is well known in the errors-in-variables literature that if there is an independent measurement error in a single independent variable, the estimated coefficient tends to be biased toward zero by a factor of proportionality called the reliability ratio (see, for example, Wayne Fuller). The reliability ratio is the ratio of the variance of the correctly measured independent variable to the sum of the variance of the correctly measured independent variable and the variance of the measurement error. We have information (in the form of estimated standard errors) on the size of the measurement error; this size varies through time, and we can assess movements in the reliability ratio through time. Reestimating Table 4 where the independent variable was a time-varying estimated reliability ratio (thereby downweighting inaccurately measured observations) did not substantially improve the significance of the results.

VIII. Conclusion

There is substantial persistence through time in rates of change in indexes of real housing prices in the cities. A change in real citywide housing prices in a given year tends to predict a change in the same direction, and one-quarter to one-half as large in magnitude, the following year. Predictable movements in real interest rates do not appear to be incorporated in prices. Our experiments with a variety of assumptions about rental rates and taxes indicate that citywide after-tax excess returns are forecastable.

While we have suggested a trading rule for individual home owners that appears to be profitable, there are still some doubts about the results. We cannot measure the dividend on housing accurately. Our measure of the dividend on housing, the BLS residential rent index, is estimated from data on rental properties which may differ in quality from owner-occupied housing, and we do not know the constant of proportionality for the index. We have given only rudimentary attention to the effects of tax laws.

There is little hope of proving definitively whether the housing market is not efficient. We see no way of obtaining an accurate historical time-series on implicit rents of owner-occupied houses. Available property tax series appear to have major deficiencies. There is not just a single income tax bracket, so any effort to model tax effects runs into definitional problems.

From the standpoint of forecasting excess returns of individual houses, such factors may be of only secondary importance anyway. The noise in individual housing prices is so great relative to the standard deviation of changes in citywide indexes that any forecastability of individual housing prices due to forecastability of citywide indexes will tend to be swamped out by the noise. Of course, this conclusion may not apply to

periods of extraordinary price movements, such as have been observed over the last few years in the northeast United States and in California.

APPENDIX
Constructing the Multiple Sales Files

The multiple sales files were constructed from several basic data sets containing large amounts of information on recorded sales of just under a million individual housing units between 1970 and 1986 (221,876 Atlanta, 397,183 Chicago, 211,638 Dallas, 121,909 San Francisco). The San Francisco data were actually drawn from the eastern part of the metropolitan area including Oakland, Berkeley, Piedmont, Hayward, and the rest of Alameda County. The data from the other three cities were drawn from the entire metropolitan areas.

The data from Atlanta, Chicago, and Dallas, as well as the data from before 1979 from San Francisco, were obtained from the Society of Real Estate Appraisers' Market Data Center in Atlanta. The data from San Francisco between 1979 and 1986 were obtained from the California Market Data Cooperative in Glendale, California, a licensee of the Society.

The data in the basic data sets were collected by members of the Society, who include many real estate agents, bank officials, and appraisers. When a transaction occurs (at the closing), members fill out a long data sheet and submit it to the Society. We have no information about how representative the membership is. We do know that the data seem to be uniformly distributed across each area and that they contain a large number of both high- and low-priced properties.

Information coded for each observation includes the exact street address, the sales price, the closing date, and the type of financing, as well as between 25 and 40 characteristics of the property depending on the city and the time period. Characteristics include numerous structural and parcel descriptors such as number of rooms, condition, lot size, and so forth. To complete our raw data set 16 separately coded files were merged.

The process of identifying repeat sales involved several steps. First, an exact match was done on the address fields. Next, properties identified as anything other than a single-family home, such as a condominium or a cooperative unit, were dropped. Third, pairs were excluded if there was evidence that the structure had been physically altered. This was done by checking the total number of rooms, the number of bedrooms, the indicated condition, and whether any rooms had been modernized.

The condition and modernization variables were recorded differently in the various data sets that had to be merged. For condition, most used ratings of excellent, good, average, fair, and poor. Because the ratings were subjective and given by different people, often many years apart, we decided to ignore small changes. Thus, a property that went from good to average was retained. Any unit that indicated a jump of two categories between sales, such as a drop from good to fair, was excluded. All properties listed in poor conditions in either period were excluded on the grounds that the rate of physical deterioration was likely to be high, and that there could well be unobservable problems reflected in price.

Whether the kitchen or a bathroom had been "modernized" was also recorded on the forms in a variety of ways. Records that indicated a modernized room were flagged, and if a flag appeared at the time of the second sale but not at the time of the first sale, the record was dropped. Of the total of 39,267 clean pairs, 57 observations appeared to be data-entry errors; the two sales prices differed by a factor close to 10.

REFERENCES

Bailey, Martin J., Muth, Richard F. and Nourse, Hugh O., "A Regression Method for Real Estate Price Index Construction," *Journal of the American Statistical Association*, December 1963, *58*, 933–42.

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

Engle, Robert F., Lilien, David M. and Watson, Mark, "A DYMIMIC Model of Housing Price Determination," *Journal of Econometrics*, June 1985, *28*, 307–26.

Fama, Eugene F., "Short-Term Interest Rates as Predictors of Inflation," *American Economic Review*, June 1975, *65*, 269–82.

Fuller, Wayne A., *Measurement Error Models*, New York: Wiley & Sons, 1987.

Gau, George W., "Public Information and Abnormal Returns in Real Estate Investment," *AREUEA Journal*, 1985, *13*, 15–31.

_____, "Weak Form Tests of the Efficiency of Real Estate Investment Markets," *Financial Review*, November 1984, *19*, 301–20.

Hansen, Lars Peter and Hodrick, Robert J., "Forward Exchange Rates as Optimal Predictors of Future Spot Rates: An Econometric Analysis," *Journal of Political Economy*, October 1980, *88*, 829–53.

Huizinga, John and Mishkin, Frederic S., "Monetary Policy Regime Shifts and the Unusual Behavior of Real Interest Rates," *Carnegie–Rochester Conference Series on Public Policy*, Spring 1986, *24*, 231–74.

Linneman, Peter, "An Empirical Test of the

Efficiency of the Housing Market," *Journal of Urban Economics*, September 1986, *20*, 140–54.

Webb, Cary, "A Probabilistic Model for Price Levels in Discontinuous Markets," in W. Eichhorn, ed., *Measurement in Economics*, Heidelberg: Physica-Verlag, 1988.

U.S. Department of Labor, Bureau of Labor Statistics, *CPI Detailed Report*, Washington: USGPO.