

Arbitrage in Housing Markets

By

Edward L. Glaeser
Harvard University and NBER

and

Joseph Gyourko
University of Pennsylvania and NBER

Draft of October 9, 2007

Abstract

Urban economists understand housing prices with a spatial equilibrium approach that assumes people must be indifferent across locations. Since the spatial no arbitrage condition is inherently imprecise, other economists have turned to other no arbitrage conditions, such as the prediction that individuals must be indifferent between owning and renting. This paper argues these non-spatial, no arbitrage conditions are difficult to use empirically for many reasons. Owned homes are extremely different from rental units and owners are quite different from renters. The unobserved costs of home owning such as maintenance are quite large, inducing added imprecision into financial no-arbitrage conditions. Furthermore, risk aversion and the high volatility of housing prices compromises short-term attempts to arbitrage by delaying home buying. Hence, it may make more sense to focus on the admittedly imprecise implications of the spatial no arbitrage condition.

This paper was prepared for the Lincoln Land Institute conference in honor of Karl (Chip) Case and his many contributions to our understanding of housing markets. Our analysis is grounded on many of his insights over the years. Finally, we appreciate the excellent research assistance of Andrew Moore on this project.

I. Introduction

Like the economic study of financial assets and wages, analysis of the housing sector rests on ‘no arbitrage’ relationships. Case and Shiller (1987, 1989, 1990) were pioneers in the study of housing price dynamics and emphasized a financial no arbitrage condition where investors earn equal risk-adjusted returns by investing in housing or other assets. Poterba (1984) and Henderson and Ioannides (1982) focus on the no arbitrage condition between renting and owning a home. Alonso (1964) and Rosen (1979) examine the implications for housing prices implied by a spatial no arbitrage condition where individuals receive similar net benefits from owning in different places.

The spatial equilibrium condition has had remarkable successes in predicting the distribution of prices and density levels within a metropolitan area (Muth, 1969) and across metropolitan areas (Roback, 1982). Yet this no arbitrage condition yields disturbingly imprecise predictions about price levels, at least by the standards of financial economics. This spatial equilibrium model clearly implies that housing should cost more in more pleasant climes, but it cannot tell us whether people are “overpaying” for California sunshine. Moreover, the heart of the model lies in spatial comparison, so it could never help us understand whether national housing prices are too high or too low.

By contrast, the other equilibrium relationships seem to offer considerably more precision. Smith and Smith (2004) and Himmelberg, Mayer and Sinai (2005) look at rent and price data and make strong inferences about whether housing prices are too high relative to the expected cost of future rents. Case and Shiller (1989) argue that the predictability of housing prices suggests excess returns for investors that run counter to the efficient markets hypothesis.

After we discuss the weaknesses of the spatial equilibrium approach in the next section in Section II, we turn to the robustness of other housing-related, no arbitrage relationships. In Section III, we first argue that it makes sense to conflate the rent-own no arbitrage relationship with the purely financial no arbitrage analysis of Case and Shiller (1989). In both cases, the key prediction of the absence of arbitrage is that there will not be excess predictable returns for owning.

Our primary conclusion is that the seeming precision of financial no arbitrage conditions in the housing market is misleading. The analogy this perspective makes between housing and stocks is troubled for several reasons. First, as we show in Section IV, the house price-to-rent ratio predicted by the no arbitrage condition is extremely sensitive to key unobserved elements of the buy-rent no arbitrage condition such as maintenance costs and expected tenure. Reasonable differences in the parameter values of these variables can generate nearly 50 percent differences in the predicted ratio between house prices and rents.

The importance of unobserved factors can be seen by simultaneously examining two financial no arbitrage conditions: a prospective investor in a house must be indifferent between becoming a landlord and investing in some other asset; and a prospective renter must be indifferent between renting and owning. As landlords have no advantage comparable to the tax shield provided by homeownership, landlords should not be willing to pay as much as an owner-occupier for the same unit of housing as a owner, at least if the landlord has the same maintenance cost and cost of capital as an owner-occupier. Reasonable parameter values suggest that itemizing owner-occupiers should be willing to pay about 40 percent more than landlords for the same property if they both

face the same costs. This gap may reflect higher maintenance costs for landlords or higher capital costs for some renters, but whatever the true explanation, any reconciliation requires that unmeasured attributes account for a 40 percent difference in predicted house price-to-rent ratios. The importance of such unobserved attributes suggests that the empirical precision of the own-rent no arbitrage equation is limited.

Section V emphasizes a second problem with using the rent-own no arbitrage condition: rental units tend to be very different from owner-occupied units. The *American Housing Survey (AHS)* documents that the vast majority of owned units are single-family detached dwellings, while rental units are highly likely to be part of a denser multifamily building. The average owner-occupied housing unit is about double the size of the typical rental unit according to the *AHS*. In addition, rental and owner-occupied units also typically are sited in different parts of the metropolitan area. Rental units tend to be closer to the urban core and are more likely to be in less attractive neighborhoods (as evaluated by residents surveyed in the *AHS*). These spatial differences may impact both the predicted level of prices and the expected level of future price appreciation.

Some researchers such as Smith and Smith (2004) have made truly heroic efforts to ensure their rental and owner-occupied properties are comparable, but this is not feasible for large scale statistical work that involves all the key markets in the country. Furthermore, given the large observable differences between rented and owned units, we suspect that unobservable differences are also considerable. Moreover, even these units are not truly comparable because of a third problem: the demand for owned units comes from a different section of the population than the demand for rental units. For example,

owner-occupiers are substantially more likely to be married and to have minor children in the home. Data from the most recent *AHS* also shows that the median income of owner households is twice that of renter households. Other sources indicate that income volatility is much greater for owners in general and for recent home buyers in particular. All this suggests that there are different demand schedules for owning and renting, which further implies that rents and prices need not be all that highly correlated over time.

Section VI turns to the problems that make it difficult to use the short term, no arbitrage relationship implied by the ability to delay purchase or sale. While there may not be many people on the margin between being a lifelong renter versus a lifelong owner, it certainly could be possible to arbitrage in the housing market by postponing a home purchase simply by remaining a renter or delaying a transition to rental status by not selling immediately. However, the ability to arbitrage by delaying the transition from rental to owning status in a declining market is limited by the high volatility of housing prices. If a buyer knows that she will have to buy, delaying the purchase creates a large amount of volatility in wealth because house prices vary so much even over annual periods. Our calibrations show that reasonable amounts of risk aversion will lead one not to delay a purchase, especially in the more expensive and volatile coastal markets.

However, risk aversion does not counterbalance the gains from delaying a sale when transitioning to rental status, largely because existing owners are likely to have much greater wealth. While homeowners looking to sell and then rent are a group that could arbitrage on the rent-own margin, according to Sinai (1997), less than four percent of owners ever transition to renting. The small and select group of people who transition from owning to renting severely limits the influence of this arbitrage channel to equalize

the returns to owning and renting. Thus, it is quite possible that substantial random shocks to housing prices will not be arbitrated away by changing the timing of a purchase or sale.

Our skepticism about the own versus rent no arbitrage condition and the no excess returns no arbitrage condition leads us back to the spatial no arbitrage condition of Alonso (1964) and Rosen (1979). This condition may not yield precise implications about price levels, but it does generate implications about the moments of housing price changes and new construction. In Section VII, we describe the results of Glaeser and Gyourko (2006) where we use the spatial no arbitrage condition to understand housing dynamics. Our results strongly support the finding of Case and Shiller (1989) that there is too much high frequency positive serial correlation in price changes. Just as Shiller (1981) finds too much variation in stock prices relative to dividends, we find that there is too much volatility in price changes relative to changes in fundamentals in the expensive coastal markets. Finally, we describe how to make more use of rental data in these exercises. Section VIII concludes.

II. The Spatial Equilibrium Model

The spatial equilibrium model requires homeowners (or renters) to be indifferent between different locations. If housing consumption is fixed, then we can write the utility function as $U(Y_i - R_i, A_i)$ where Y_i represents income, R_i is the cost of housing and A_i represents a vector of j location-specific amenities. The term $Y_i - R_i$ represents cash after housing costs, and we are assuming that non-housing prices are constant across

space. The spatial equilibrium assumption implies that $U(Y_i - R_i, A_i)$ is constant across space or:

$$(1) \quad dR_i = dY_i + \frac{1}{U_1} \sum_{j=1}^J \frac{\partial U}{\partial A_j} dA_{i,j}$$

where $A_{i,j}$ denote the different elements in the vector of amenities. Differences in housing prices across space are associated either with higher income levels or higher amenity levels. The spatial equilibrium assumption allows us to treat one area within the U.S. as a reservation locale, and we denote its income as \underline{Y} , its housing prices with \underline{R} , and its amenity levels with \underline{A}_j for each amenity j . We then use a first-order Taylor approximation to find that:

$$(2) \quad R_i \approx \underline{R} + (Y_i - \underline{Y}) + \frac{1}{U_1} \sum_{j=1}^J \frac{\partial U}{\partial A_j} (A_{i,j} - \underline{A}_j)$$

In each location, the housing cost is approximately equal to the housing cost in the reservation locale plus the difference in income between location i and the reservation locale plus the sum of all of the amenity differences times the marginal utility of each amenity divided by the marginal utility from income.

While this equation implies a tight, even one-for-one, relationship between the changes in the flow of housing costs and housing prices, it does not directly tell us about the level of prices. It might be possible to use this to look at rent differences over space, but for reasons that we will discuss later, we think that renters and rental units are sufficiently unrepresentative of a metropolitan area that we are skeptical about using rents in this fashion. If we want to use this equation to deal with prices, we need to make further assumptions that relate housing prices with per period housing costs.

Following Poterba (1984) and others, the per period cost of housing can be written $(1 - \tau)(r + p)H(t) - [H(t + 1) - H(t)]$, where $H(t)$ denotes housing prices at

time t , τ denotes the income tax rate, r denotes the interest rate and p denotes the local property tax rate. If we make the heroic assumption that we are in a steady state where housing prices are expected to be constant over time, then the per period housing costs are just $(1 - \tau)(r + p)H(t)$, or $\mu H(t)$, where μ denotes a fixed ratio between housing prices and housing costs or $(1 - \tau)(r + p)$. If housing prices were known to appreciate at a fixed rate α then the value of μ is $(1 - \tau)(r + p) - \alpha$.

If we know the value of μ , then the model makes a hard quantitative prediction about the relationship between changes in income and house prices. Specifically, every dollar increase in income should be associated with a $\frac{1}{\mu}$ increase in housing costs. The relationship between housing costs and incomes across metropolitan areas in 2000 is shown in Figure 1. While the slope is undeniably positive, the coefficient is 5.6 (standard error of 0.99), meaning that a \$1 increase in income is associated with a \$5.60 increase in housing prices.¹ This would be compatible with the model if μ was equal to 0.18. This number is higher than standard user cost estimates which range from 7.5-12 percent. Such user costs would suggest that the coefficient on income should lie between 8 and 12, yet we generally find that it is far lower.²

We can still save the spatial equilibrium model by appealing to omitted variables. For example, higher income places might also have lower amenity values, especially if the higher income levels are compensating for lower amenities as in Rosen (1979). Alternatively, higher prices might not accurately reflect different housing costs because

¹ The underlying data are from our 2006 working paper which uses information on 116 metropolitan areas for which we have consistent price and income data over more than two decades. The house price data are for the median quality home from the 1980 census, with the house value in 2000 reflecting the appreciation in the OFHEO repeat sales index for each metropolitan area. Median family income is from the 2000 decennial census. All values are in \$2000.

² Different cross sections and data generate different results, of course. If we use 1990 data, the coefficient estimate increases to 6.2, but that still implies higher user costs than most researchers believe are sensible.

we are ignoring any heterogeneity in expected housing cost appreciation. Thus, the spatial equilibrium model is salvageable, but any claims about its tight precision are not. The one numerically precise implication that comes out of the model doesn't seem to fit the data, and if the model is correct, then unobserved variables must be quite important.

In addition, predictions about prices and amenities are never particularly tight. Certainly, the model predicts that prices should rise with positive amenities, as indeed they do. Figure 2 shows the positive connection between housing prices and median January temperature across the same sample of metropolitan areas in 1990. However, there is no external estimate of the value of $\frac{1}{U_1} \frac{\partial U}{\partial A_j}$ that would enable us to know whether the observed relation of a \$1,158 higher house price (standard error of \$549) for each extra degree of winter warmth is too high or too low. Indeed, housing price regressions of this kind are generally used to provide such estimates since nothing else is available.

By first differencing the linear approximation to the spatial no arbitrage relationship, we also gain predictions about the dynamics of housing prices:

$$(3) R_{i,t+1} - R_{i,t} \approx \Delta_t + Y_{i,t+1} - Y_{i,t} + \frac{1}{U_1} \sum_{j=1}^J \frac{\partial U}{\partial A_j} (A_{i,j,t+1} - A_{i,j,t})$$

where $\Delta_t = \underline{R}_{t+1} - \underline{R}_t - (\underline{Y}_{t+1} - \underline{Y}_{j,t}) - \frac{1}{U_1} \sum_{j=1}^J \frac{\partial U}{\partial A_j} (\underline{A}_{j,t+1} - \underline{A}_{j,t})$. This equation implies that changes in housing costs should be tightly connected to changes in income and changes in amenities.

While this equation implies a one-for-one relationship between the changes in the flow of housing costs and housing prices, it does not directly tell us about the level of prices either. Moreover, assuming that we are in a steady state where housing prices are fixed is logically inconsistent with a regression that is examining heterogeneity in

housing price changes. If we want to use this equation to deal with prices, we need to make further assumptions that relate housing prices with per period housing costs. In particular, we need to make assumptions about the extent to which housing price changes are expected or unexpected.

At one extreme, we can assume that any shocks to income or amenities are completely unexpected. In that case, the model predicts that a \$1 increase in income will continue to be associated with a $\frac{1}{\mu}$ dollar increase in housing prices. This assumption is surely counterfactual since local income changes are quite predictable (Glaeser and Gyourko, 2006).

The other extreme is to assume that local income changes are entirely known in advance. To create simple closed form solutions, we can go so far as to assume amenities and housing costs in the reservation locale are constant over time and that the gap in income between location i and the reservation locale is growing by $g_{i,Y}$ dollars per period. In this case,

$$(4) H_{i,t} \approx \frac{1}{\mu} \left(\underline{H} + (Y_{i,t} - \underline{Y}_t) + \frac{1}{u_i} \sum_{j=1}^J \frac{\partial U}{\partial A_j} (A_{i,j} - \underline{A}_j) \right) + \frac{g_{i,Y}}{\mu^2}$$

and $H_{i,t+1} - H_{i,t} \approx \frac{(Y_{i,t+1} - Y_{i,t}) - (Y_{t+1} - Y_{j,t})}{\mu}$. Expected income changes will have exactly the same impact on housing price changes as unexpected income changes, as long as those income changes are part of a long-run trend in income appreciation.

An intermediate option that yields a slightly different result is to assume that there is a one-time increase in income between time t and $t+1$ that is anticipated, but that there will not be any more shocks to income after that point. In that case, the impact of

housing changes is much smaller: $H_{i,t+1} - H_{i,t} \approx \frac{(Y_{i,t+1} - Y_{i,t}) - (Y_{t+1} - Y_{j,t})}{1+\mu}$. Price changes

will exist, and they will be predictable, but they will be much smaller than either in the case where price changes are unexpected or where they reflect a long-run trend. Since we know little about the information that people have about income shocks, the model does not deliver a tight implied relationship between housing price changes and income changes. Instead the implied coefficient could range from $\frac{1}{\mu}$, which could be more than ten, to $\frac{1}{1+\mu}$, which is less than one.

Figure 3 plots the actual relationship between house price changes and income changes across metropolitan areas between 1980 and 2000 for the same 116 metropolitan area sample used above. House prices again use the 1980 Census median value as the base value, with the relevant OFHEO metropolitan area price index used to scale prices over time. The change in income is the 20 year difference in median family income between the 1980 and 2000 censuses. As expected, the figure shows a robust positive relationship, with a coefficient of 5.1 (standard error of 0.42) from a simple regression of 20-year price changes on 20-year income changes. Happily, this number lies between 1 and 10, so it does not reject the spatial equilibrium model. However, the bounds implied by the model are so loose that it would have been shocking for the model to be rejected.

In this section, we have shown that the spatial no arbitrage condition fails to yield tight predictions about the relationship between housing prices and income, which we treat as the “fundamental” in that model. This weakness surely plays a role in explaining why real estate economists have been attracted to other no arbitrage relations, which we consider next.

III. The Arbitrage of Buying and Renting

Case and Shiller's (1989) pioneering work on housing price dynamics discusses both the no arbitrage conditions for buyers and for investors. The no arbitrage condition for buyers is usually a no arbitrage condition between buying and renting, which may either involve a lifetime indifference or an indifference between buying (or selling) and renting for a short time period. Case and Shiller (1989) themselves emphasize the decision of a buyer to purchase today or to wait for a year. Those authors also discuss the possible decision of a buyer who is looking at whether or not to increase housing consumption. In this case, estimating the costs of delay must include an estimate of the inconvenience associated with consuming too little housing. Since that inconvenience level surely is impossible to directly measure, this approach cannot offer much precision, and we are not surprised that subsequent work has focused primarily on the owner-renter no arbitrage condition. We will first focus on that condition, and then turn to the investors' no arbitrage condition.

The simplest version of the financial approach to housing involves a one-period indifference condition where consumers receive the same return from owning or renting the identical housing unit. We will later emphasize that risk aversion is likely to be far more important in the context of housing markets than it is in financial assets, but we begin with a representative risk neutral individual who is considering buying a house at time t and leaving the city with probability one at time $t+1$. The buyer must pay property taxes of p times the housing price H . The interest rate at which this person can both borrow and lend is given by r . Both property taxes and interest payments are deductible for owner-occupiers. If we further assume this person earns $Y(t)$ dollars and faces a

marginal tax rate of τ , then the owner's user cost of housing will equal $(1-\tau)(r+p)H(t) - [H(t+1) - H(t)]$. If the same individual rents, housing costs equal $R(t)$.

Poterba (1984) and others have emphasized additional costs of housing, too. For example, the average owner spends nearly \$2,100 per year on maintenance, although there is substantial measurement error in this variable (Gyourko and Tracy, 2006). In addition, this observable component of maintenance misses the time and effort that owner-occupiers put into caring for their homes. The economic depreciation of a house is also difficult to measure. Of course, ownership may also bring with it hidden benefits such as the ability to customize the housing unit to one's own needs. We let $\delta H(t)$ denote the net unobserved costs of being an owner-occupier, or one's own landlord as it were.

With these costs, indifference between owning and renting for a risk-neutral resident implies that $R(t) = [(1-\tau)(r+p) + \delta]H(t) - E[H(t+1) - H(t)]$, where the final term represents the expected capital gain on the housing unit. Iterating this difference equation and imposing a transversality condition on housing prices yields the familiar formula that prices are the appropriately discounted sum of rents: $H(t) = \sum_{j=0}^{\infty} \frac{R(t+j)}{(1 + (1-\tau)(r+p) + \delta)^{j+1}}$.

As this equation shows, it is impossible to determine the appropriate price of housing as a function of rents without knowing the long term path of rents. Since expected future rents are certainly unobservable, this in turn creates ambiguity in the formula. One approach is to assume that rents will rise at a constant rate a , so that $R(t+j) = (1+a)^j R(t)$. If so, then equation (5) holds

$$(5) H(t) = \frac{R(t)}{(1-\tau)(r+p) + \delta - a}.$$

This provides a precise prediction for house value if rents, true maintenance and depreciation, interest rates, property tax rates, marginal income tax rates, and expected capital gains are known, and if unobserved influences on user costs are small in magnitude.³ Drawing on Himmelberg, Mayer and Sinai (2005) for some baseline numbers, if $\tau=0.25$, $r=0.055$, $p=0.015$, $\delta=0.025$, and $\alpha=0.038$, the nominal price-to-rent ratio is 25.⁴ While many of these variables can be observed in some manner, we think there is a reasonable amount of uncertainty about what true maintenance, expected appreciation, and even what the relevant interest and tax rates are. For example, if expected appreciation actually is one percentage point higher, the multiple increases about one-third to 34, and larger changes can be generated by incorporating relatively minor adjustments to the other parameters. Thus, any reasonable sensitivity analysis is going to result in a fairly wide bound for what prices ‘should be’ in a given market.⁵

³ It is noteworthy that Poterba (1984), who generally is credited with introducing this model into mainstream economics, neither considered the own-rent margin nor equated the utility flow from owning with the observed rental price of a house. He used the user cost formula to determine the cost to owners, which then shifts the demand for the quantity of housing.

⁴ This is only slightly higher than the ratio predicted by the more complex formula used by those authors.

⁵ We have greater faith in the value of the comparative statics suggested by equation (5) than in its ability to justify the level of prices. However, there is considerable debate in the literature over one important result involving the impact of interest rates on house prices. Equation (5) suggests a powerful relationship between interest rates and house prices, and McCarthy and Peach (2004) and Himmelberg, Mayer and Sinai (2005) have relied on it to justify currently high house prices at least partially as a function of historically low interest rates. In contrast, Shiller (2005, 2006) argues that there is no economically or statistically significant relationship between house prices and interest rates over any reasonably long period of time. When we regressed the real value of the median quality home from 1980 (using the OFHEO index as described above) on the real 7-year interest rate using data from the last 30 years, the results indicated that a one percentage point increase in interest rates is associated with only a 2-3 percent rise in house prices. The $R^2=0.12$, which is well below the nearly 2/3rds of variation in house prices that can be accounted for by metropolitan area fixed effects. However, this is a very complex issue that cannot be definitively answered within the confines of such a simple static model. For example, one can imagine a dynamic setting in which interest rates mean revert and in which homeowners either can refinance loans or expect to sell and buy another home within a few years. In either case, temporary rises in rates need not lead to substantially higher debt service costs (in present value terms) that are capitalized into lower house values if refinancing costs are low and borrowers believe rates will drop in the relatively near future. In general, mean reversion in interest rates implies that we should see far less connection between current rates and house prices than is predicted by the constant interest rate version of the model (Glaeser and Gyourko, 2006). Essentially, this points up another problem in using a static user cost approach to deal with issues that may be inherently dynamic in nature.

IV. Differences in the Owner-Occupied and Rental Stocks

One underappreciated problem with using the rent-own no arbitrage condition to make inferences about housing prices is that rental units are generally quite different from owner-occupied housing. In this section, we document some of the key differences. We also show that renters and owners are very different people, which suggests that the forces driving demand for rental and owner-occupied housing are also likely to be different.

We begin by documenting a number of physical characteristics of owner-occupied and rental units in Table 1. For this analysis, we rely primarily on the latest *American Housing Survey (AHS)* from the year 2005. Perhaps the most striking fact about renting and owning is the very strong correlation between unit type and physical structure. The 2005 *AHS* shows that 64.3 percent of owner-occupied housing units were of the single-family, detached unit type, while only 17.7 percent of rental units were of that type. The vast majority of rental units are in multiple-unit buildings, not single-unit, detached dwellings.

Naturally, these types of units are of very different sizes. Figure 4, which is taken from Glaeser and Gyourko (2007), plots the median square footage of owned versus rented units using data from the last twenty years of the *AHS*. The median owner-occupied unit is nearly double the size of the median rented housing unit in the United States. Per person consumption of space also varies widely by tenure status. Housing consumption per capita among owner-occupied households is now over 700 square feet, while that for renters is about 450 square feet (Glaeser and Gyourko, 2007).

Not only is the owner-occupied versus rental stock physically quite different, they tend to be located in different parts of the metropolitan area, as well as in different quality neighborhoods. The suburban dominance of owner-occupancy is highlighted in Table 1 by the fact that less than one-third of all owned units were in the central cities of metropolitan areas according to the 2005 *AHS*. Ownership has become more widespread in America's central cities, but nearly half of all rental units still are located in cities (row 2, column 2 of Table 1). Owner-occupied units tend to be in better neighborhoods, too. The *AHS* asks its survey responders to rate their neighborhoods on a scale of 1-10. Just looking at those who gave their neighborhoods very high scores of 9 or 10 score shows that almost one-half of owners believe they live in the highest quality areas, while only one-third of renters felt the same way (row 3, Table 1).

Just as owned units are different from rented units, owner-occupiers are quite different from renters. Perhaps most importantly, owners are substantially richer. The median nominal income of owner-occupier households was \$53,953 versus \$24,651 for renter households according to the 2005 *AHS* (row 4, Table 1). Household types also tend to differ systematically by tenure status, as indicated by the fact that the probability an owner-occupier household is a married couple with minor children present is nearly double that of a renter household (bottom row of Table 1).

Other data indicates that the volatility of the two income series also is very different. For example, comparing the incomes over time of recent buyers in the *Home Mortgage Disclosure Act (HMDA)* data with that for the mean in an area as reported by the *Bureau of Labor Statistics (BLS)* finds the volatility of recent buyer income roughly double that of the average income in the same market (Glaeser and Gyourko, 2006). A

similar pattern can be seen specifically for the New York City market in the *New York City Housing and Vacancy Survey (NYCHVS)* data. A simple regression of the income of recent buyers (defined as those who bought within the past two years) on BEA-reported per capita real income for that market finds that recent buyer income goes up by \$1.29 for every \$1.00 increase in BEA-measured income. Moreover, the same source reveals that renter-household incomes are less volatile than average. They increase by only \$0.47 for every \$1.00 rise in per capita income in the city.⁶

Taken literally, all this indicates that the variance of income shocks for renters is only a small fraction of that for owners or for the general population. If so, rent series should be more stable than house prices. On the aggregate level, Leamer (2002) has emphasized that house prices have grown much more quickly than rents. In the 44 markets for which we have both consistent rent data from a prominent industry consultant and constant quality repeat sales indices, Table 2 documents that the annual appreciation rate for housing is 1.9 percent since 1980, while that for rents is only 0.5 percent.⁷ Table

⁶ Because the *NYCHVS* provides much smaller samples, the regression results are based on averages of individual respondents over two-year windows. Effectively, there are only nine observations after averaging, and while the regression coefficients are statistically significant, one clearly does not want to make too much of this. The underlying regression results are as follows. For owners,

$$\text{Recent Buyer Real Income} = -32,451 + 1.29 * \text{BEA Per Capita Real Income} \\ (15,102) \quad (0.19)$$

There are nine observations (one for each survey year), the $R^2=0.87$, and standard errors are in parentheses. For renters,

$$\text{Renter Real Income} = -2,885 + 0.47 * \text{BEA Per Capita Real Income} \\ (5,436) \quad (0.07)$$

The number of observations again is nine, the R^2 still is 0.87, and standard errors are in parentheses.

⁷ The rental series is from REIS, Inc. The company does not report a constant quality series, but their data are consistently measured in the sense they reflect the answers to a question about asking rents on higher quality apartment complexes in major U.S. markets. Rent data are very rare and there is little existing analysis of the robustness of such series. We found that the REIS asking rent series is strongly positively correlated with the rent subindex of the local CPI index that the Commerce Department computes for about

2 also reports results for a handful of representative major markets in which price growth typically is at least double that of rent growth. Similar patterns with relatively low rent volatility also exist if one breaks the data into different time periods.

One explanation for the mismatch in the growth of housing prices and rents is that housing prices represent the cost of accessing higher quality housing units, while rental prices represent the cost of accessing lower quality units. Rising incomes and rising income inequality could easily mean that demand has increased more for higher quality units. Gyourko, Mayer and Sinai (2005) argue that housing prices have risen more steadily for metropolitan areas with higher amenity levels.

Of course, an empirical mismatch between house price growth and rent growth still could be explained by a purely financial model if other factors such as interest rates or expected house price appreciation themselves are changing. We have already noted the debate about the role of interest rates, so that remains an unsettled issue. There also is not much convincing evidence that the differences between home prices and rents are positively correlated with price appreciation. A proper user cost model implies that the user costs of housing minus rents should equal expected house price appreciation.

In Table 3, we report the results from regressing actual house price appreciation on the gap between user costs and rents, using the user cost data from Himmelfarb, Mayer and Sinai (2005).⁸ Over one, three and five year horizons, there is a negative, not

25 areas nationwide. REIS also reports an 'effective rent' series that allegedly reflects discounts or premiums being charged tenants. That series is not positively, and sometimes is negatively, correlated with the local CPI rent subindex numbers. Hence, we do not use it in any of the analysis reported here.

⁸ We thank Todd Sinai for providing their underlying data. Because we need user costs before expected housing appreciation, we added back their appreciation component, which is based on the long-run average annual real appreciation rate over 1940-2000 in Gyourko, Mayer and Sinai (2005). We then create a shorter-run expected price change variable by multiplying the user costs before appreciation figure by the real value of a 1980 quality home and then subtracting off real asking rents. The house price variable is

a positive, relationship between actual house price appreciation and the change in house prices forecast by a user cost model. While Shiller-type animal spirits certainly could be behind this, our point simply is that there is no strong evidence that variation in the relationship between prices and rents is systematically related to accurate assessments of house price appreciation.

This leads us to conclude that the house price and rent series reflect the costs of two different types of housing. The differences seem so large that it probably is best to think of them as reflecting different demands for two related, but not directly comparable, markets. Of course, there still will be some sort of indifference relationship between owned and rental housing, but quantifying this relationship in the way suggested by the standard user cost approach will tend to be empirically misleading. The indifference relationship appears to be sufficiently weak that there is abundant opportunity for the financial costs of owning to diverge significantly from the financial costs of renting.

V. The Importance of Omitted Costs

Smith and Smith (2005) represent the best effort that has been made to deal with the often stark differences in rental versus owner-occupied units. However, their approach still faces the problem that owners and renters are likely to be quite different people. Moreover, this work also needs to deal with the challenge that unobserved factors in equation (5), such as maintenance costs, may be very important and could lead to quite different predictions about the appropriate relationship between housing prices and rents.

computed by scaling the mean house value in each market as reported in the 1980 census by the OFHEO repeat sales index appreciation for each year. The rent data are from REIS, Inc, and are discussed above.

Both theory and data suggest that unobserved influences on user costs are likely to be large in magnitude. A rental property must involve two different agents—the renter and the landlord—and both of them have relevant no arbitrage conditions. The renter must be indifferent between renting and owning. The landlord must be indifferent between owning a housing unit and renting it out and investing his capital in something else. This no arbitrage condition implies a second way of evaluating the appropriate price of housing, but the price implied by the investor’s no arbitrage condition will be very different from the price implied by the renter’s no arbitrage condition, unless omitted variables are quite important.

To illustrate this, we assume that the investor also has the ability to borrow at interest rate r . The relevant no arbitrage condition is that the net present value of revenues from the property is zero. Gross revenues equal the rent received each period, or $R(t)$. Property taxes on the unit are the same as for an owner-occupier. However, we allow for net maintenance costs to differ, so that they equal $\delta_I H(t)$ for the investor. Profits are taxed, but the tax rate is irrelevant if there are zero profits. Hence, the zero profit condition is given by equation (6),

$$(6) R(t) + E[H(t+1) - H(t)] - (r + p + \delta_I)H(t) = 0.$$

The same zero profit condition holds if the investor either can lend money at rate r or buy a house, with all revenues being taxed at a rate τ_I . The relevant indifference condition is given by $(1-\tau_I)rH(t) = (1-\tau_I)[R(t) - (p + \delta_I)H(t) - E\{H(t+1)-H(t)\}]$, which again yields equation (6). Thus, the tax rate on the investor does not impact the relationship between prices and rents in this simple model.

Iterating equation (6) and imposing a transversality condition implies that

$$H(t) = \sum_{j=0}^{\infty} \frac{R(t+j)}{(1+r+p+\delta_I)^{j+1}}. \text{ Again presuming that rents rise at a constant rate } a \text{ so that}$$

$$R(t+j) = (1+a)^j (R(t)), \text{ the formula implies}$$

$$(5') \quad H(t) = \frac{R(t)}{r+p+\delta_I-a}.$$

There are two ways to use equations (5) and (5'). First, we can assume that $\delta=\delta_I$ and ask how much bigger the price-to-(net) rent ratio should be for owner-occupiers than for landlords. For owner-occupiers, the no arbitrage condition predicts a price-to-net rent

$$\text{ratio of } \frac{1}{(1-\tau)(r+p)+\delta-a}. \text{ For the investor-landlord, the price-to-net rent}$$

$$\text{relationship predicted by the no arbitrage condition is } \frac{1}{r+p+\delta-a}. \text{ The two}$$

relationships are the same only when the owner-occupier does not deduct interest and taxes.

However, the housing literature that uses the rent-own no arbitrage decision to deduce housing prices generally assumes that owners are deducting interest.⁹ Recall from above that using Himmelberg, Mayer and Sinai's (2005) assumptions for our

$$\text{parameter values finds the price-to-rent ratio given by } \frac{1}{(1-\tau)(r+p)+\delta-a} \text{ is about 25}$$

⁹ Data on itemization by tenure status is not directly reported by the IRS, but it is only natural to presume that homeowners are more likely to itemize. Nationally, only 35.7 percent of all tax returns filed in 2005 did so. Given the nearly 69 percent homeownership rate estimated for that year, at least half of owners did not itemize, even if we assume that all itemizers own their home. However, itemization rates are higher in higher house price areas, which is consistent with more owners in those markets being able to deduct local property taxes and mortgage interest payments. For example, 39.9 percent of California returns, 38.8 percent of New York returns, and 45.2 percent of New Jersey returns itemized in 2005.

(~1/.0395). However, the investor-landlord's no arbitrage condition implies a price-rent ratio of $\frac{1}{r + p + \delta - a}$, which equals 17.5. This means that owner-occupiers should be willing to pay about 45 percent more for the same house than should a landlord (25/17.5 ~1.45).

One way of interpreting this is that if we think that price-to-rent ratio eliminates arbitrage between renting and owning, then housing is nearly 50 percent too expensive to eliminate the arbitrage between being a landlord and other forms of investment. This gap increases in higher appreciation or inflation environments because the tax subsidy to owner occupiers rises with inflation (Poterba, 1984). If expected inflation increases so that the nominal interest rate rises to 8 percent and the rate of appreciation equals .063

instead of .038, then the ratio $\frac{1}{(1 - \tau)(r + p) + \delta - a}$ rises to over 30, while the ratio

$\frac{1}{r + p + \delta - a}$ remains at 17.5. Indeed, it is relatively easy to envision environments in

which the price-to-rent ratio implied by the owner's no arbitrage condition literally would be double that implied by the renter's no arbitrage condition.

There are many possible ways that we can reconcile these seemingly incompatible predictions about price-to-rent ratios. We wrote the model so that different maintenance/depreciation rates could do the job. The two no arbitrage conditions will imply the same price-to-rent ratio when $\tau(r+p) = \delta - \delta_i$, or when the difference in the maintenance rates just equals the difference in the tax advantage provided owner occupiers. If $\tau=0.025$, and $r+p=0.07$, then this would mean that the maintenance costs are 0.0175 higher for the owner-occupier than for the landlord. Only if it costs more for the

owner to keep up his home can we explain why landlords would ever buy and rent at the same prices that make owner-occupiers indifferent between owning and renting.

A second way to reconcile the two no arbitrage conditions is that the landlord's cost of capital might be lower than the owner-occupier's cost of capital. Perhaps the marginal buyer has more difficulty making a down payment or negotiating the loan process. If maintenance costs were the same, then landlords would need to face interest rate costs that were 175 basis points lower than prospective tenants in our simple example.

Alternatively, risk tolerance might differ between owners and landlords. Perhaps the marginal buyer has a relatively short time horizon in the city and does not want to face the risk of housing price shocks (Sinai and Souleles, 2005), while landlords are diversified and remain immune to those shocks. There are many unobserved factors that could explain the seeming incompatibility of the two no arbitrage conditions.

Our point is that unobservable elements must be quite important in housing markets because they need to explain a 40+ percent difference in the price-to-rent ratios predicted by the landlord's no arbitrage condition and the owner-occupier's no arbitrage condition. The magnitude of these unobserved factors should make us wary of believing these conditions can be used to definitively answer whether prices are too high or too low. After all, reasonable parameters suggest that a price level that looks entirely appropriate from the perspective of an owner-occupier simultaneously looks 45 percent too high from the perspective of a potential investor-landlord.

VI. Risk Aversion and the One Period No Arbitrage Condition

We now turn to the one-period no arbitrage condition between owning and renting. While owners and renters are generally quite different people, individuals are often both renters and owners over the course of their lifetimes. When they transition from renting to owning, or the reverse, individuals have the opportunity to delay purchase, or sale, to exploit predictability in housing prices. Case and Shiller (1989) specifically focus on the ability to exploit excess returns by delaying consumption for one period.

In this section, we argue that the ability to exploit any predictable excess returns is compromised by the interaction between risk aversion and the volatility of housing prices. While individuals may be effectively risk neutral with respect to one individual stock that represents a small share of their overall portfolio, housing usually is the dominant asset for most homeowners. Normal year-to-year variation in housing prices can create significant swings in an individual's total wealth. The magnitude of these swings creates an incentive for anyone who knows that they are going to buy next year to buy today and for anyone who knows that they are going to sell next year to sell today. Thus, there appear to be even more limits to arbitrage in the housing market than there are in the financial markets (Shleifer, 1999).

Consider the case of a household that knows with certainty that it eventually will own a home in a given market, and assume that it can either buy at time t or wait until time $t+1$. To simplify the notation from above, we abstract from local property taxes and assume away any unobserved costs associated with maintenance or other aspects of

owning. The only two flow costs remaining are debt service, where the interest rate still is denoted r , and known maintenance and depreciation, which is denoted as M .

We assume that this household is maximizing its expected wealth, denoted $E(V(Wealth_{t+1}))$, where $Wealth_{t+1}$ refers to wealth net of housing costs as of time $t+1$. By assumption, the household must have bought a home by that date. If the household buys at time t , its wealth at time $t+1$ is predictable. The household's total welfare will equal $V((1 + (1 - \tau)r)(Y - H(t)) - M(t))$. If it rents at time t and then buys, its wealth at time $t+1$ will be stochastic and will equal

$$E(V((1 + (1 - \tau)r)Y - R(t) - H(t+1))).$$

To calibrate the model, we will use a second-order Taylor series expansion for the function $V(\cdot)$ and assume that $H(t+1) = H(t) + \overline{H(t+1) - H(t)} + \varepsilon(t)$, where $\varepsilon(t)$ is mean zero and $\overline{H(t+1) - H(t)}$ is the predictable component of the change in housing prices. With these assumptions, delay only makes sense if:

$$(7) \quad (1 - \tau)rH(t) + M(t) - R(t) - \overline{H(t+1) - H(t)} > \frac{\sigma Var(\varepsilon)}{(2 - z\sigma)((1 + (1 - \tau)r)(Y - H(t)) - M(t))}$$

where σ denotes the coefficient of relative risk aversion, i.e.

$$\sigma = -\frac{((1 + (1 - \tau)r)(Y - H(t)) - M(t))V''((1 + (1 - \tau)r)(Y - H(t)) - M(t))}{V'((1 + (1 - \tau)r)(Y - H(t)) - M(t))}, \text{ and } z \text{ represents}$$

the ratio of expected one period gains from delaying to total wealth if the individual does

$$\text{not delay, i.e. } z = \frac{(1 - \tau)rH(t) - \overline{H(t+1) - H(t)} - R(t) + M(t)}{((1 + (1 - \tau)r)(Y - H(t)) - M(t))}.$$

Equation (7) provides a useful bound for the plausible amount of expected losses that would justify waiting one year given reasonable values of risk aversion. The standard deviation of annual housing price changes in our sample of 116 metropolitan areas is just over \$9,100. If the coefficient of relative risk aversion is 2 and if we assume non-home wealth of \$50,000 for a person buying at time t , then the expected gains from waiting would need to be at least \$1,750.¹⁰ Thus, risk aversion causes the plausible gulf between the user costs of owning and rental housing costs to increase by nearly \$150 per month, even for a renter household with \$50,000 in non-housing wealth.

To help gauge whether the potential benefit of exploiting short-run predictability can counter this risk aversion affect, we begin by regressing the one-year, forward-looking change in house prices on observables such as the current house price and macroeconomic variables such as the long-term real rate and real gross domestic product. We also include metropolitan area dummies, so knowledge of average, one-year price changes also is presumed. All the observables are statistically significant predictors of the coming year's price change. Table 4 reports the distribution of predicted one-year changes.¹¹ Just under one-third of the expected one-year changes in house prices are negative (31 percent to be precise), and only 18 percent of the cases involve expected losses of more than -\$1,750, which is required to generate positive returns to a renter household delaying purchase for one year, given our assumptions.

This calculation, however, assumes that the variation in housing prices is constant across markets, which obviously is not the case. Hence, in the second column of the

¹⁰ To simplify the calculation, this result also assumes that $(1-\tau)rH(t) + M(t) - R(t) = 0$, not only that non-housing wealth or $((1+(1-\tau)r)(Y - H(t)) - M(t)) = \$50,000$.

¹¹ The precise equation estimated is $P_{i,t+1} - P_{i,t} = \alpha + \beta * P_{i,t} + \gamma * 10yrRealRate_t + \delta * RealGDP_t + \eta * MSA_i + \epsilon_{i,t}$. All results are available upon request.

table, we report the distribution of the total gains from delay for our hypothetical household using information on price volatility at the metropolitan area level. Predicted price changes still are estimated via the specification with metropolitan area fixed effects, lagged house price, and the other economic variables. However, the variance of ε is computed separately for each metropolitan area by using the relevant residuals from the equation used to predict housing price changes. Once again assuming that

$(1 - \tau)rH(t) + M(t) - R(t)$ equals zero, $((1 + (1 - \tau)r)(Y - H(t)) - M(t))$ equals \$50,000 and σ equals two, the formula for expected gains minus risk aversion-related losses then

$$\text{equals } -\frac{H(t+1) - H(t)}{100,000 + 2H(t+1) - H(t)} - \frac{2\text{Var}(\varepsilon)}{100,000 + 2H(t+1) - H(t)}.$$

The second column in Table 4 reports on the distribution of net benefits from our hypothetical renter household delaying purchase of a home for a year. They are positive in only 26 percent of the cases, and a look at the results by metropolitan areas indicates that it is in the high price volatility coastal markets where risk aversion almost always more than counterbalances the gross benefits of waiting to purchase in a declining market. As indicated by the results in column one, house prices are expected to rise in most markets in most years. However, even in the highest appreciation markets in the northeast region and coastal California, our naïve forecasting equation does generate expected declines in the early 1980s and the early 1990s when general economic conditions were quite poor. Nevertheless, in no case does our simple calculation show a positive return to delaying purchase in any of the five major coastal California markets in our sample (Los Angeles, San Diego, San Francisco, San Jose, and Santa Barbara). For the Boston and New York City areas, the return to delay is positive only once—in 1980

when forecasted price declines were large enough to outweigh the costs associated with risk aversion.

The reason is the very high volatilities of price changes in these markets. The values of $Var(\varepsilon)$ among these seven large coastal markets, range from a low of \$175 million in Boston to a high of \$572 million in San Francisco. In contrast, the impact of risk aversion is much less in many interior markets. For example, Atlanta's $Var(\varepsilon)$ value is only \$12.9 million. Its home prices were expected to fall in only 8 of the 26 years for which we can forecast, but in each of those years the return to our hypothetical renter household delaying purchase for year is positive.

In sum, this arbitrage opportunity only has value if price declines can be expected, and that is not the normal condition in our housing markets. However, even if we reasonably can expect price declines over the coming year in markets such as New York, Boston, and the Bay Area, the volatility of their house price changes is more than enough so that risk aversion eliminates any gain from delaying the purchase of a home. Hence, it seems unlikely that renters considering changing tenure status in these markets will find this potential arbitrage opportunity to be of value.

Of course, there also is the possibility to arbitrage renting and owning among those individuals who are moving from owner-occupied to rental housing. In this case, people could delay a year in order to take advantage of a rising market. To consider this issue more formally, we continue to assume that households maximize $E(V(Wealth_{t+1}))$. If a household sells immediately, its expected wealth is deterministic and expected welfare will equal $V((1 + (1 - \tau)r)(Y + H(t)) - R(t))$. If the household waits a year, then its time $t+1$ wealth is stochastic and expected welfare will equal

$E(V((1 + (1 - \tau)r)Y + H(t + 1) - M(t)))$). Again using a second-order Taylor series expansion, we see that it is sensible to wait if and only if:

$$(8) \quad R(t) - (1 - \tau)rH(t) - M(t) + \overline{H(t + 1) - H(t)} > \frac{\sigma \text{Var}(\varepsilon)}{(2 - z_2 \sigma)((1 + (1 - \tau)r)(Y + H(t)) - M(t))},$$

where σ continues to denote the coefficient of relative risk aversion, and z_2 represents the ratio of expected one period gains from delaying to total wealth if the individual does

not delay, i.e. $z_2 = \frac{R(t) - (1 - \tau)rH(t) - M(t) + \overline{H(t + 1) - H(t)}}{((1 + (1 - \tau)r)(Y + H(t)) - M(t))}$.

The impact of risk aversion should be smaller here because wealth should be much larger. To show this more clearly, we now calculate the distribution of gains from waiting a year to sell a home, again computing $\text{Var}(\varepsilon)$ at the metropolitan area level. As before, we assume that $(1 - \tau)rH(t) + M(t) - R(t)$ equals zero, and σ equals two, but now we presume that $((1 + (1 - \tau)r)(Y + H(t)) - M(t))$ equals \$250,000. With these assumptions the expected risk-adjusted gain from waiting a year can be written as

$$\overline{H(t + 1) - H(t)} - \frac{2\text{Var}(\varepsilon)}{500,000 - 2\overline{H(t + 1) - H(t)}}.$$

The third column of Table 4 reports our estimates of the distribution of expected gains from an existing owner delaying sale. While it often did not make sense to delay a purchase decision to take advantage of falling prices, especially in the more volatile markets, there generally are substantial gains from delaying a sales decision to take advantage of rising prices. Over 70 percent of the observations exhibit positive returns to this potential arbitrage opportunity. Not only are prices expected to appreciate in most cases, but the much larger assumed wealth substantially mitigates the impact of risk aversion so that it rarely counterbalances the benefits of delay even in the most volatile

markets. In principle, the population of home-owners looking to sell and rent represents one group that really could arbitrage along the own-rent margin whenever prices are expected to increase.

However, there is a reason to expect that the impact of this arbitrage possibility on housing prices is quite small—namely, very few people actually transition from owning to renting. Sinai (1997) documents that transitions from owner-occupancy to renter status are quite rare. Working with a 1970-1992 panel of observations from the *Panel Study of Income Dynamics*, he shows that less than four percent of owners ever engage in such a tenure transition, and of those that do, about one-third transition back to ownership within two years.¹² Hence, we are skeptical that this group can be a real force for creating an equilibrium where renting and owning returns are equalized. This only reinforces our view that the relevant indifference relationships between owning and renting are not as tight as a purely financial perspective might indicate.

V. Using Price and Rent Data Together to Understand Housing Markets

We have just detailed a number of reasons why it is extremely difficult to use the buy-rent no arbitrage condition to produce precise predictions about housing prices. That said, we do believe that there is much to be learned from the use of rents and prices together to understand housing dynamics. In this section, we discuss three ways in which these data can be employed to add insight into housing prices.

¹² Capital gains taxation rules explain the short tenure spells in this case. A household must trade up in value within two years to be able to rollover any gains from the original sale. Our point is not about the arcana of the tax code, but to illustrate that a large fraction of the transitions from owner-occupancy to rental status are very short term and probably not related to the arbitrage we are discussing. In addition, Sinai (1997) reports that falls in income have an especially large impact on the probability of this type of tenure transition (see his Table 4), which suggests that households making this move are suffering some type of negative income shock, not trying to arbitrage along the rent-own margin.

The first use of rents lies in prediction without theory. Rents may add predictive power to housing price change regressions even if we are not sure why they have this predictive power. Table 5 details the results from nine regressions where changes in housing prices have been regressed on initial characteristics. The basic specification is:

$$(9) \text{Log} \left(\frac{\text{Price}_{t+j}}{\text{Price}_t} \right) = \beta_1 \text{Log}(\text{Price}_t) + \text{MSA Dummies} + \text{Year Dummies}.$$

We repeat this specification for j equal to 1, 3, and 5 years, using the same house price variable described above. As Case and Shiller (1989) first showed us, there is much that is predictable about house price changes, simply from knowing previous price levels.

We then repeat this basic specification using the price-to-rent ratio instead of prices themselves, as in equation (9')

$$(9') \text{Log} \left(\frac{\text{Price}_{t+j}}{\text{Price}_t} \right) = \beta_2 \text{Log} \left(\frac{\text{Price}_t}{\text{Rent}_t} \right) + \text{MSA Dummies} + \text{Year Dummies}.$$

Note that using the price-to-rent ratio is associated with uniformly higher t-statistics, as well as a higher R² for the 1- and 3-year price change horizons. Over longer five year periods, one cannot reject the null that the elasticity of price changes with respect to the price-to-rent ratio is -1.

The final specification reported in Table 5 enters prices and rents separately, i.e.:

$$(9'') \text{Log} \left(\frac{\text{Price}_{t+j}}{\text{Price}_t} \right) = \beta_3 \text{Log}(\text{Price}_t) + \beta_4 \text{Log}(\text{Rent}_t) + \text{MSA Dummies} + \text{Year Dummies}.$$

Note that both prices and rents are highly significant at standard confidence levels for high and low frequency price changes. Higher levels of rents tend to predict higher price growth holding prices constant. And, the R²'s are uniformly higher than in the base case that includes only prices (row 1).

These regressions show that incorporating both prices and rents does improve our ability to explain housing price changes over time. As such, brute empiricism can clearly benefit from including both prices and rents. However, as our discussion above should have made clear, there are real problems in deciphering the meaning of these results.

A particularly naïve view might be that the negative correlation between the price-to-rent ratios and future price growth seems to reject the view given in equation (5) that $\frac{H(t)}{R(t)}$ equals one divided by $(1 - \tau)(r + p) + \delta - \alpha$. If that were the case, then higher housing price-to-rent ratios should predict future appreciation, not future depreciation. Clearly, they do not and especially not over longer time intervals.

One interpretation of these results is that the market is fundamentally irrational and that prices don't internalize reasonable expectations of future housing price growth, but instead reflect some kind of irrational exuberance (Shiller, 2005). An alternative interpretation is that rents are telling us about a related, but different, market than owner-occupied housing. Rents are, by and large, reflecting the cost of housing in lower quality homes in the inner city. House values are, by and large, reflecting the cost of housing in the suburbs. These lower quality inner city homes are not a perfect substitute for the suburban homes, but they are at least something of a substitute. If higher rents are associated with higher housing price appreciation, this might reflect the fact that rents are giving us new information about the state of the region's economy that are not fully embedded in house prices. Higher rents might well mean that demand is robust not only for high end housing, but for low end housing, and this could easily mean that the future of region is brighter. According to this view, the role of rents in the housing price

regression does not reflect irrationality, but rather the natural role of providing more information about the future of the region's economic strength.

While rents can naively be inserted into a regression aimed at maximizing predictive power, it is harder to actually connect housing prices and rents with a structural model to test its implications. Glaeser and Gyourko (2006) write down a straightforward model of housing dynamics and then test its implications using housing prices and permits, but we do not look significantly at rents for reasons discussed more fully below. In the model, high frequency changes in demand for housing are driven by changing economic conditions within a region. We use the model to generate predictions about the moments of price and quantity fluctuations.

That model shows that unobservable differences in the information structure will have enormous impact on the predicted high frequency correlations between prices and new construction. If people recognize economic shocks only when they occur then the predicted correlation between price changes and new construction will be almost perfect. If people learn economic shocks a period before they occur, then the predicted correlation between price change and new construction will be almost zero. As outside researchers have little ability to assess the actual information that people have, we believe that these results mean that it makes little sense to look at high frequency correlations between prices and construction.

For the same reason, it makes little sense to look at high frequency correlations between rents and housing prices. Again, if people discover economic shocks only when they occur, then the correlation between price changes and rent changes will be extremely high. If people learn about future changes a period or more ahead of time, then this

correlation will be significantly lower. Results like these make us wary about focusing on the correlations between price innovations and rents.

However, the model does deliver predictions that are more robust to changes in the information structure. For example, the actual variances of price change and new construction are implied by the variation in underlying economic shocks, and these relationships are not particularly sensitive to the timing of new information. Our empirical work suggests that the variability of prices for the median market in the U.S. seems close to the variability predicted by the model. However, the variability of prices for the most expensive models is far too high to be explained by the underlying economic variation. This excess variation is the housing price analogue of Shiller's (1981) finding of excess variation in the stock market.

The model also predicts the autocorrelations of both price and quantity changes. Notably, despite the fact that the model has no irrationality, there is every reason to expect that price changes will be predictable. In the long run, prices are predicted to mean revert both because economic shocks appear to mean revert and because new construction becomes available. In fact, the model predicts a level of mean reversion over five years that is almost identical to the level of mean reversion that we see in the data.

The model is less successful in predicting the high frequency positive serial correlation that is also a feature of the data. The high frequency positive serial correlation in the OFHEO data is probably biased upwards because it contains appraisal data and because of inaccuracy in the timing of sales. However, using much better sales data purged of these problems, Case and Shiller (1989) also documented substantial price

persistence at high frequencies. This serial correlation is not predicted by our spatial equilibrium model, and we agree with the original Case and Shiller conclusion that this momentum provides a challenge to conventional models of housing price dynamics.

Glaeser and Gyourko (2006) do little with rents for two reasons. There is good reason to believe that observed rent levels may indeed understate the true volatility of rents because of long-term explicit and implicit-relationship between landlords and tenants. This problem becomes even more severe in areas with rent control. Moreover, the observed data on median or average incomes may well tell us about the marginal homebuyer, but it does not tell us about the marginal renter. As noted above, data from the *New York City Housing and Vacancy Survey* indicate that the variability of renter income is less than one-quarter of the variability of owner income. This lower variability should predict low variation in rents.

How then could rents be brought into a model of housing dynamics that started with a spatial no arbitrage assumption? The first requirement is to have good high frequency income data for a set of metropolitan areas that reflected the income of potential renters. For some larger metropolitan areas, this could potentially be done with the *American Community Survey*, but it would be difficult to get a significant sample of metropolitan areas. The second requirement would be to get high frequency data on new rental contracts, preferably involving new tenants. Such data would presumably be free of any longer-term commitments between tenants and landlords.

These tasks are not easy, but they offer some promise of enabling us to use rental data to test the predictions of the spatial no arbitrage model. While we recognize the

difficulties of these tasks, other approaches, like focusing on a no arbitrage condition between owning and renting seem even less promising.

VI. Summary and Conclusion

Economics forms predictions about housing prices with no arbitrage conditions, and different researchers have emphasized different ways in which housing prices can be arbitrated. The traditional urban approach has been to emphasize the absence of arbitrage across space, but this approach never delivers too much precision. A more financial approach has emphasized the ability of investors to arbitrage prices and the ability of owners to arbitrage between owning and renting.

The major point of this paper is that the seeming empirical precision of these more financial approaches is illusory. While the conceptual ability to arbitrage between owning and renting is clear, the ability to use this insight empirically is limited. Owned units and rented units are extremely different. Unobserved components of housing costs, like maintenance, are quite large. Owners and renters are quite different people, and risk aversion creates a substantial cost to delaying a purchase especially. For these reasons, we are skeptical that rental data can tell us much about the appropriate price of a house.

Instead, we believe that the spatial equilibrium framework offers a more promising approach for understanding the nature of housing markets. Our past work in this area suggests that some seeming anomalies of housing markets, like the high mean reversion of prices over five year intervals, is quite compatible with a rational spatial equilibrium model. Other seeming anomalies, like high frequency positive serial correlation of price changes and high volatility in coastal markets, seem to be much

harder to reconcile with such a market, just as Case and Shiller (1989) have suggested. It would be possible to bring rents into such a model if we had better data on the income series of potential renters, and if we had better data on new rental contracts. We hope future work will follow this path.

References

- Alonso, William. *Location and Land Use*. Cambridge, MA: Harvard University Press, 1964.
- Case, Karl and Robert Shiller. "Prices of Single Family Homes Since 1970: New Indexes for Four Cities", *New England Economic Review*, September/October 1987: 45-57.
- _____. "The Efficiency of the Market for Single-Family Homes", *American Economic Review*, Vol. 79, no. 1 (March 1989): 125-137.
- _____. "Forecasting Prices and Excess Returns in the Housing Market", *Journal of Real Estate Finance and Economics*, Vol. 18, no. 4 (1990):
- Glaeser, Edward and Joseph Gyourko. "Housing Dynamics", National Bureau of Economic Research Working Paper No. 12787, December 2006.
- _____. "Housing Policy and House Prices: How to Make Housing More (Not Less) Affordable", mimeo, American Enterprise Institute, August 2007.
- Gyourko, Joseph, Chris Mayer and Todd Sinai. "Superstar Cities", National Bureau of Economic Research Working Paper No. 12355, June 2006.
- Gyourko, Joseph and Joseph Tracy. "Using Home Maintenance and Repairs to Smooth Variable Earnings", *Review of Economics and Statistics*, Vol. 88, no. 4 (November 2006): 736-747.
- Henderson, Vernon and Yannis Ioannides. "A Model of Housing Tenure Choice", *American Economic Review*, Vol 73, no. 1 (1983): 93-113.
- Himmelberg, Charles, Chris Mayer and Todd Sinai. "Assessing High House Prices: Bubbles, Fundamentals and Misperceptions", *Journal of Economic Perspectives*, Vol. 19, no. 4 (Fall 2005): 67-92.
- Leamer, Edward. "Bubble Trouble: Your Home Has a P/E Ratio, Too", mimeo, UCLA Anderson School of Management, June 2002.
- McCarthy, Jonathan and Richard Peach. "Are Home Prices the Next Bubble?" *Economic Policy Review*, Vol. 10, no. 3 (December 2004): 1-17.
- Muth, Richard. *Cities and Housing. The Spatial Pattern of Urban Residential Land Use*. University of Chicago Press: Chicago, IL 1969.

Poterba, James. "Tax Subsidies to Owner-Occupied Housing: An Asset Market Approach", *Quarterly Journal of Economics*, Vol. 99, no. 4 (November 1984): 729-745.

Roback, Jennifer (1982). "Wages, Rents, and the Quality of Life", *Journal of Political Economy*, Vol. 90, no. 4 (December 1982): 1257-78.

Rosen, Sherwin (1979). "Wage-Based Indexes of Urban Quality of Life". In *Current Issues in Urban Economics*, edited by Peter Mieszkowski and Mahlon Straszheim. Baltimore: Johns Hopkins University Press, 1979.

Shiller, Robert. "Do Stock Prices Move Too Much to be Justified by Subsequent Changes in Dividends", *American Economic Review*, Vol. 71, no. 3 (1981): 421-436.

_____. (2005). *Irrational Exuberance*. Princeton University Press: Princeton, NJ 2006 (2nd edition).

_____. "Long-Term Perspectives on the Current Housing Boom", *Economists' Voice*, Vol. 3, no. 4 (March 2006): 1-11.

Shleifer, Andrei and Robert Vishny. "The Limits of Arbitrage", *Journal of Finance*, Vol. LII, no. 1 (March 1997): 35-55.

Sinai, Todd. "Taxation, User Cost and Household Mobility Decisions", mimeo, The Wharton School, University of Pennsylvania, December 1997.

Sinai, Todd and Nicholas Souleles. "Owner-Occupied Housing As a Hedge Against Rent Risk", *Quarterly Journal of Economics*, Vol. 120, no. 2 (May 2005): 763-789.

Table 1: Comparing the Owner-Occupied and Rental Housing Stocks^a		
	<i>Owner-Occupied Housing</i>	<i>Renter-Occupied Housing</i>
%Single-Family Detached Unit Type	64.3	17.7
%Located in Central Cities	30.5	45.7
%Rating Their Neighborhoods as Excellent ^b	45.6	34.2
Median Household Income in 2005	\$53,953	\$24,651
%Married Households with Minor Children	27.6	15.4

Notes:

^aData are from the 2005 *American Housing Survey* unless otherwise noted.

^bWe label as neighborhood as excellent if the survey respondents gave it a rating of 9 or 10 on a 1-10 scale.

Table 2: Comparing House Price and Rental Growth 44 Markets with Continuous Rent Data from REIS, Inc.		
	1980-2006	
	Average Annual Rent Growth	Average Annual Price Growth
44 Markets	0.51%	1.88%
San Francisco	1.96%	3.93%
Boston	2.06%	4.37%
Los Angeles	1.29%	3.62%
Atlanta	0.22%	1.06%
Chicago	0.83%	2.20%
Phoenix	-0.20%	2.19%

Notes: The rent data are from REIS, Inc. House price appreciation rates are computed from the OFHEO price index. All data are in real terms.

Table 3: Is Actual Real House Price Appreciation Consistent with Forecasts from a User Cost Model?

$P_{i,t+n} - P_{i,t} = \alpha + \beta*(F_{i,t+1} - F_{i,t}) + \delta*Year_t + \eta*MSA_i + \varepsilon_{i,t}$ <p>where $P_{i,t+n} - P_t$ = change in real house prices in market I, and $F_{i,t+1} - F_{i,t}$ = one-period change in real house prices forecast by user cost model</p>	
1-year horizon ($P_{i,t+1} - P_t$):	$\beta = -0.81$; $n=1119$, $R^2 = 0.40$, cluster by msa (0.22)
3-year horizon ($P_{i,t+3} - P_t$):	$\beta = -9.60$; $n=358$, $R^2 = 0.57$, cluster by msa (0.83)
5-year horizon ($P_{i,t+5} - P_t$):	$\beta = -14.03$; $n=224$, $R^2 = 0.64$, cluster by msa (1.14)

Notes: Data on user costs were provided by Todd Sinai and are identical to that used in Sinai, Mayer, and Himmelberg (2005). See their paper for the details behind the user cost calculation. See our footnote #8 for details on the calculation of the user cost model forecast of real house price changes.

Table 4: Estimating the Benefits of Short-Term Predictability			
	<i>Distribution of 1-yr Price Changes^a</i> ($P_{i,t+1} - P_{i,t}$)	<i>Distribution of Net Gains from Delaying Purchase^b</i> $-\overline{H(t+1) - H(t)} - \frac{2\overline{Var(\varepsilon)}}{100,000 + 2\overline{H(t+1) - H(t)}}$	<i>Distribution of Net Gains from Delaying Sale^c</i> $\overline{H(t+1) - H(t)} - \frac{2\overline{Var(\varepsilon)}}{500,000 - 2\overline{H(t+1) - H(t)}}$
10 th Percentile	-\$2,698	-\$15,352	-\$2,864
25 th Percentile	-\$612	-\$8,089	-\$775
50 th Percentile	\$2,361	-\$3,199	\$2,144
75 th Percentile	\$6,163	\$112	\$5,609
90 th Percentile	\$10,802	\$2,179	\$9,739

Notes:

- a.) The underlying specification estimated regresses the one-year, forward-looking change in house prices on a series of observables as follows: $P_{i,t+1} - P_{i,t} = \alpha + \beta * P_{i,t} + \gamma * 7yrRealRate_t + \delta * RealGDP_t + \eta * MSA_i + \varepsilon_{i,t}$, where $P_{i,t}$ reflects house price in metropolitan area i in year t , $7yrRealRate$ is the real interest rate on 7-year Treasuries (calculated as in Himmelberg, Mayer and Sinai (2005)), $RealGDP$ is real gross domestic product from the Economic Report of the President, MSA_i is a vector of metropolitan area dummies, and ε is the standard error term.
- b.) Net gain from delaying purchase for one year for a renter household with \$50,000 in non-housing wealth and a relative risk aversion coefficient equal to 2. See the discussion in the text for more detail
- c.) Net gain from delaying sale for one year for an owner household with \$250,000 in wealth and a relative risk aversion coefficient equal to 2. See the discussion in the text for more detail.

Table 5: Price Changes Within Market Over Time Dependent Variable: $\text{Log}(P_{i,t+j}/P_{i,t})$ i=metropolitan area i, t=year, $j=1, 3, 5$		
(9) $\text{Log}\left(\frac{Price_{t+j}}{Price_t}\right) = \beta_1 \text{Log}(Price_t) + \text{MSA Dummies} + \text{Year Dummies}$		
1-year horizon, $j=1$ $\beta_1 = -0.034$ (0.013), $R^2=0.42$, $n=1144$	3-year horizon, $j=3$ $\beta_1 = -0.391$ (0.047), $R^2=0.53$, $n=1056$	5-year horizon, $j=5$ $\beta_1 = -0.864$ (0.079), $R^2=0.71$, $n=968$
(9') $\text{Log}\left(\frac{Price_{t+j}}{Price_t}\right) = \beta_2 \text{Log}\left(\frac{Price_t}{Rent_t}\right) + \text{MSA Dummies} + \text{Year Dummies}$		
1-year horizon, $j=1$ $\beta_2 = -0.110$ (0.015), $R^2=0.46$, $n=1144$	3-year horizon, $j=3$ $\beta_2 = -0.570$ (0.052), $R^2=0.53$, $n=1056$	5-year horizon, $j=5$ $\beta_2 = -0.984$ (0.073), $R^2=0.70$, $n=968$
(9'') $\text{Log}\left(\frac{Price_{t+j}}{Price_t}\right) = \beta_3 \text{Log}(Price_t) + \beta_4 \text{Log}(Rent_t) + \text{MSA Dummies} + \text{Year Dummies}$		
1-year horizon, $j=1$ $\beta_3 = -0.101$ (0.015), $\beta_4 = 0.199$ (0.039) $R^2=0.48$, $n=1144$	3-year horizon, $j=3$ $\beta_3 = -0.570$ (0.052), $\beta_4 = 0.533$ (0.085) $R^2=0.59$, $n=1056$	5-year horizon, $j=5$ $\beta_3 = -0.984$ (0.073), $\beta_4 = 0.578$ (0.096) $R^2=0.70$, $n=968$

Notes: Standard errors in parentheses. Specifications estimated on 44 metropolitan areas with both OFHEO house price and REIS apartment rent data for the 1980-2006 time period.

Figure 2: House Prices and Winter Warmth Across Metropolitan Areas, 1990

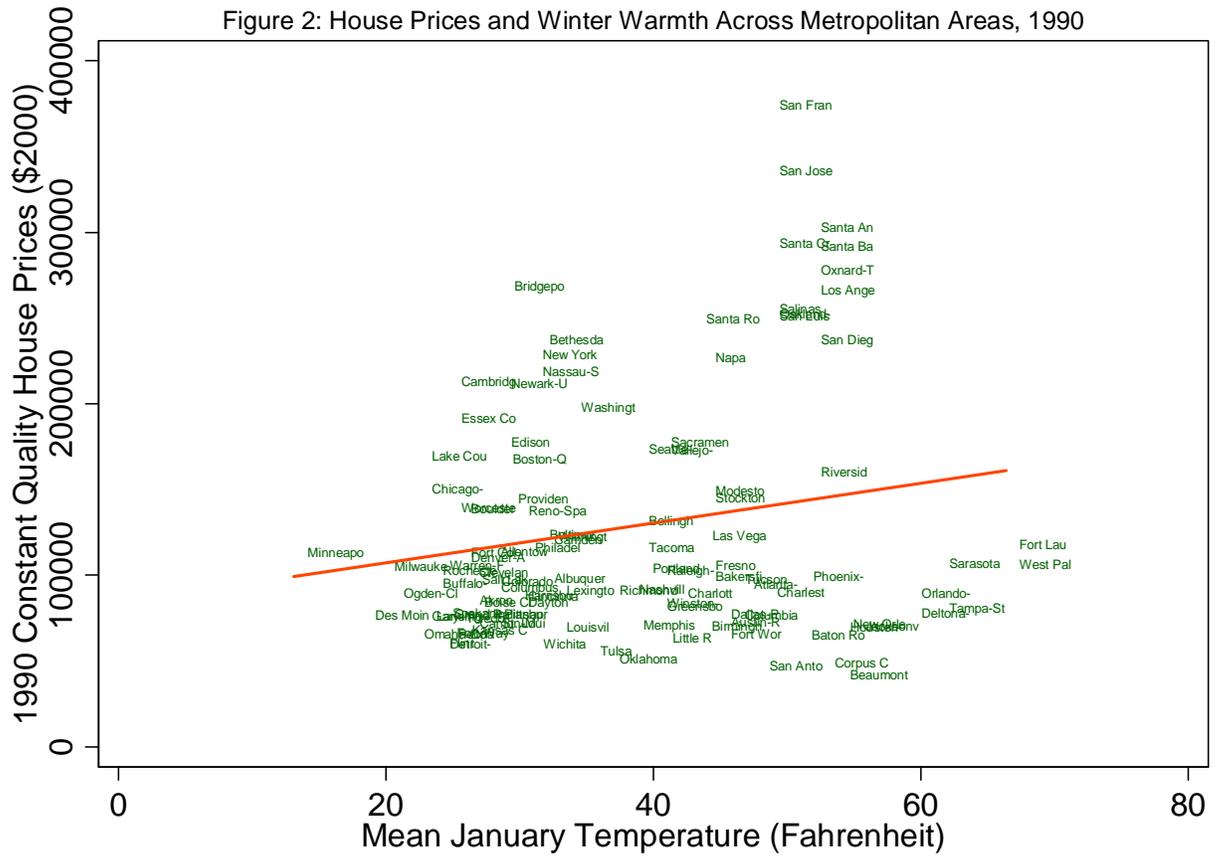


Figure 3: 20-Year Changes in House Prices and Incomes, 1980-2000

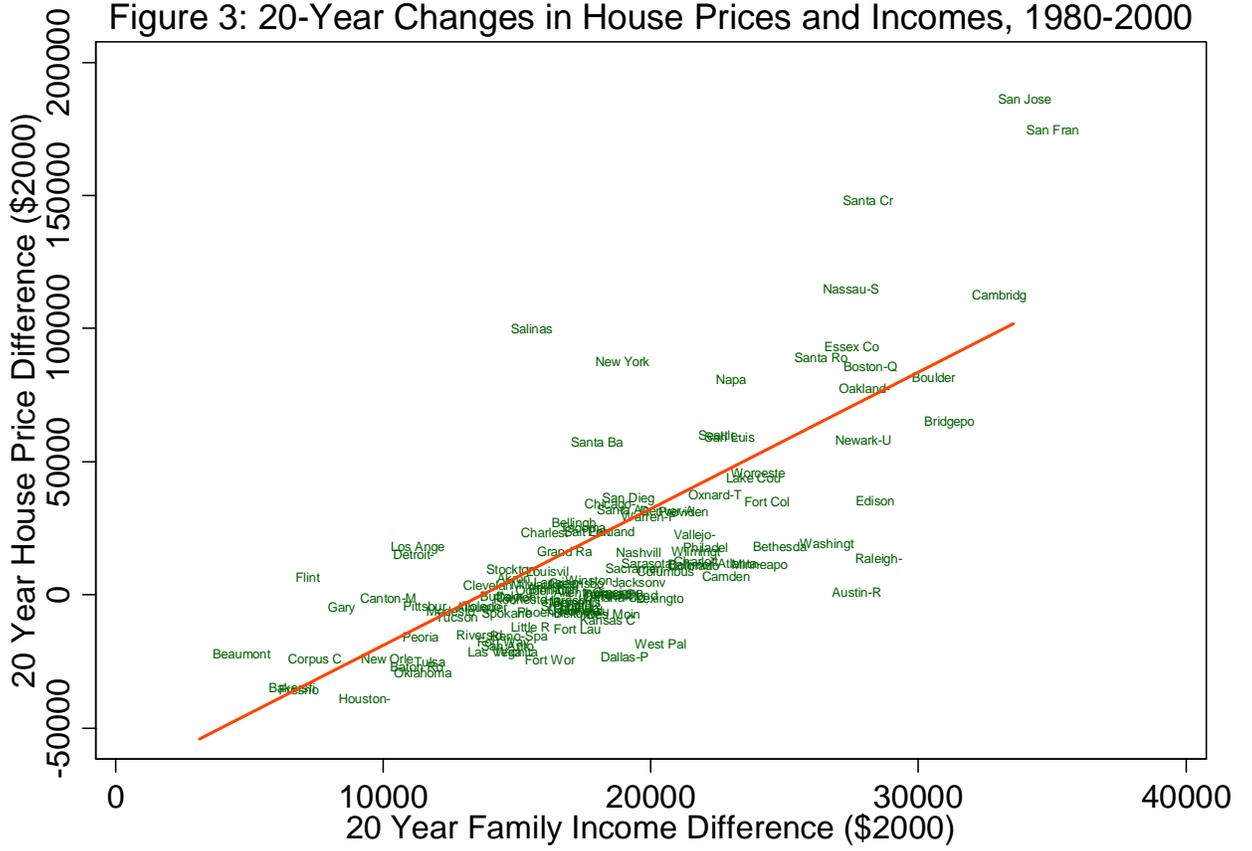


Figure 4: Housing Unit Size: Owner-Occupied vs. Rental Units

