

The puzzling divergence of aggregate rents and user costs, 1978-2001

Randal Verbrugge
Division of Price and Index Number Research (3105)
Bureau of Labor Statistics
2 Massachusetts Ave NE
Washington, DC 20212
Email: verbrugge.randal@bls.gov

First version: June 18, 2003
This version: June 24, 2004

Abstract

User costs are typically taken to be the appropriate measure of the cost of owning durable goods. In the standard frictionless theory, user costs of homeownership equal rents. It is well-known that ex-post measures of homeowner user costs are much more volatile than rents. I demonstrate here that the divergence between user costs and rents extends well beyond this difference. Even ex-ante measures of user costs are four orders of magnitude more volatile than rents; and even more surprisingly, over a two-decade period the series do not even appear to share a common trend. Furthermore, rent changes are almost unrelated to user-cost driving forces. This leaves statistical agencies in a quandary if they wish to measure homeowner costs: user-cost changes are so volatile that, if used directly, they would completely dominate movements in aggregate price indices; and changes in user costs are not approximated by (far less volatile) rent changes.

Rough Draft: Please do not cite or quote or circulate; comments welcome.

Acknowledgements: Thanks are due to the HOUSAD group at the BLS, and to Paul Emrath, Josh Gallin, and Elaine Maag; however, they are not responsible for remaining errors, which are intended as a test of the reader's knowledge. The views here are those of the author, not the BLS or other BLS members.

1. Introduction

In computing consumer price indices, the US Bureau of Labor Statistics (BLS) – and several other statistical agencies in the US and abroad – use the rental equivalence method to estimate changes in homeownership costs. In this method, rent inflation in neighboring rental markets – possibly adjusted for costs of utilities and so on – is taken to be an accurate measure of inflation in homeownership costs.

User costs are generally thought to be the appropriate theoretical tool to use in understanding the costs of owning durable goods such as houses. In the standard frictionless theory, rent levels should equal homeowner user costs – i.e., what one's home could command in rent should equal the *ex ante* (or expected) user cost. Thus, market rents have been viewed as an appropriate *measurement tool* for user costs. (And one might argue that hypothetical rents are the correct measure of the flow of services consumed by a homeowner.)

Under the presence of differential tax treatments (such as the ability to deduct mortgage interest payments) and transactions costs (such as realtor fees and moving expenses), user costs become highly idiosyncratic, and rents and user costs diverge. As of yet, it is unknown whether average rent *changes* track average user cost *changes*, at least for some appropriate average of idiosyncratic user costs.

Thus, on theoretical grounds it is uncertain that rent changes are an adequate proxy for user cost changes.

The purpose of this note is not to contribute to this theoretical issue. Instead, it is to point out the highly divergent behavior of rents and (estimates of) user costs (and between rent changes and changes in user costs). Simply put, rents and estimated user costs can diverge markedly from each other for long periods of time – an empirical finding which suggests that rents are *not* a good measure of user costs (and highlights a puzzle about how rents are actually set). Furthermore, the volatility of ex-ante user cost changes is fully four orders of magnitude greater than that of rent changes, so that statistical agencies evidently do not have the option of using direct estimates of user costs as an alternative measure of the costs of homeownership.

2. Data and construction of the detached-home rent index

Data used are the aggregate OFHEO US house price index (which is intended to be an index tracking price changes of constant-quality homes), the average contract rate on commitments for 30-year conventional fixed-rate first mortgages in the US, the aggregate US rent index from the CPI, confidential post-1987 US CPI data on rents of detached housing units (as described below), and the US Census rental vacancy rate. Data are all quarterly.

A main component of an estimate of user costs is home prices. The most widely-used home price data series is the OFHEO house price index. This index is constructed using the repeat sales method, as in Case and Shiller (1987, 1989), and is described in Calhoun (1996). This method limits the extent to which changes in the composition of the sample can influence the estimated index – since only price *changes* on the same property are used in estimating the index. As Gallin (2003) points out, the OFHEO index is not a quality-adjusted price in that it does not control for changes such as improvements (or deterioration).

The goal of this paper is to directly compare estimated user costs to rents. To do this, one must construct a measure of user costs that is comparable to the rental data. Thus, at a minimum, the type of quality adjustment on the rent series should be similar to that done on the user cost series, and the type of structures for which rents are obtained should be similar to those for which user costs are estimated. Along both of these dimensions of comparability, the CPI rent index is not an entirely appropriate comparison to an OFHEO-index-based user cost series. The CPI rent index is constructed using rents of many types of rental properties (such as apartments) whose prices do not enter the OFHEO home price index. Furthermore, the CPI rent index includes rents of some rent-controlled apartments. In addition, it does not treat utilities symmetrically with owner-occupied housing. (The CPI's rental-equivalence index controls for latter two problematic issues, but the formula used to compute it changed markedly over the period in question, making it inappropriate for the comparisons attempted below.) Finally, while the CPI rent index does track rents in the identical dwelling for extended periods of time, it also, to some extent, adjusts for quality changes in the dwellings.

To address these data comparability issues (and to develop a rent index that is more directly comparable to the OFHEO house price index), I constructed a monthly rent index using confidential post-1987 CPI rent data comprised only on rents on single detached dwellings, with

no quality adjustment. In theory, then, the dwellings whose rents are tracked by this index should be similar to the dwellings whose prices are tracked by the OFHEO price index. Furthermore, it is possible to use methods akin to the OFHEO repeat-sales methodology to construct a monthly rent index, making the constructed rent index as close as possible to the OFHEO price index.

To explain how I construct the rent index, I must first discuss some aspects of BLS rent collection. Periodically, the BLS selects a new sample for rent-collection in a particular city, and divides this sample into six subsamples ($i = 1, \dots, 6$), each subsample corresponding to one of the first six months of the year. In January (for example), rent data is collected on the January sample; rent data is not subsequently collected on any unit in the January sample until July; rents are again collected on each unit in the January sample the following January. The temporal pattern of rent collection continues for every unit in the January sample, until either an entirely new sample is drawn, or the unit in question drops out of the sample. Only six-month rent changes on the same unit are used to construct the rent index.

Since sample i units never appear in sample j , it is impossible to construct a single rent index using methods similar to that used in constructing the OFHEO price index. (Put differently, there is no information linking the level of the index in February to that in January; there is only information linking the level of the index in February to that in the previous August.) In principle, however, similar assumptions and methodology could be used to construct *six* rent indices – one corresponding to each sample – which should all share the common national market rent trend over the period.

However, to date attempts to adapt this method to the CPI rent data have proved unsatisfactory. Instead, I construct six detached rent index series using a related method. It is fairly standard practice in the housing research literature to characterize individual house prices as arising from a stochastic process in which the average rate of change is represented by a market index, and the dispersion and volatility of individual house prices around this market index are modeled as log-normal diffusion processes. This assumption forms the basis of the repeat-sales method. The corresponding assumption for rents is that the log of the rent of an individual home i at time t equals a market rent $B_{i,t}$ plus an idiosyncratic Gaussian random walk $G_{i,t}$ plus an idiosyncratic white noise component $N_{i,t}$, as in

$$\ln(\text{rent}_{it}) = B_{i,t} + H_{i,t} + N_{i,t}$$

This implies that

$$\begin{aligned}\ln(\text{rent}_{i,t+6}) - \ln(\text{rent}_{i,t}) &= B_{i,t+6} - B_{i,t} + H_{i,t+6} - H_{i,t} + N_{i,t+6} - N_{i,t} \\ &= B_{i,t+6} - B_{i,t} + v_{i,t+6}\end{aligned}$$

where $v_{i,t}$ is white noise. Under this assumption, the national average of 6-month rent changes for sample j should be a reasonable estimate of the change in the national market rent over that period;¹ and thus, the average 6-month change in sample j may be used to move index j forward that six months. (As in the OFHEO repeat sales method, this limits the extent to which changes in the composition of the sample can influence the estimated index.)

The six series are then merged into a single series by selecting the initial index number for each series such that the sum of squared changes in the resulting combined series is as small as possible. As is apparent in Figure 4 below, after transforming this detached-home monthly index into the quarterly frequency, it appears to be quite similar to the CPI rent index – although the detached-home index features a slightly lower average inflation rate over the period for which it may be constructed.

Finally, a key component in a user cost series is the interest rate. The mortgage interest rate series used here implies that the user cost in question is the user cost appropriate to a *marginal* (new) rather than average consumer. (This is more problematic during periods when the mortgage interest rate is rising, since existing homeowners always have the option of refinancing when mortgage interest rates drop.) The rent index, in contrast, is explicitly an average consumer index: it tracks actual rents being paid by the population.

3. User costs

In principle, the ex ante user cost is simply the expected annual cost associated with purchasing a house, using it for one year, and selling it at the end of the year.² Given the presence of risk-neutral landlords, and under the assumption of no transactions costs, this should equal the market rent for an identical home.

¹ The sample size each month is on the order of 1000 observations.

² In this note, transactions costs and financing constraints (such as minimum down payments) are ignored. The relationship between market rents and user costs which take such complications into consideration has yet to be fully explored. (However, see Diaz and Luengo-Prado (2004) and Martin (2004).) The standard frictionless theory, which implies that rents equal user costs, is explicated in Gillingham (1980, 1983) and Dougherty and Van Order (1982).

Three different (annual) user cost formulas are employed. The first ignores the preferential tax treatment given to homeowners:

$$usercost_t = P_t (i_t + \gamma - E\pi_t^h) \quad (1)$$

where P_t is the price of the home; i_t is the mortgage interest rate; γ is the sum of depreciation, maintenance, and property tax rates (and potentially a risk premium) – all assumed constant; π_t^h is the 4-quarter home price appreciation between now and 1 year from now; and E represents the expectation operator (thus, the final term is expected annual home price appreciation). A 1-year appreciation rate is used since we are considering the *annual* user cost, in order to remain comparable to the typical rental contract. This formula makes no distinction between the different interest rates faced by equity and debt; this assumption is relatively common in the literature.

The final two user cost measures are constructed from a second user cost formula which explicitly takes account of the preferential tax treatment given to homeowners (and assumes that the homeowner is itemizing):

$$usercost_t = P_t \left(i_t (1 - \tau_t^{Fed}) + \tau_t^{prop} (1 - \tau_t^{Fed}) + \tilde{\gamma} - E\pi_t^h \right) \quad (2)$$

Here, τ_t^{prop} is the property tax rate, τ_t^{Fed} is the marginal federal income tax rate, and $\tilde{\gamma}$ includes depreciation and maintenance costs (and perhaps a constant risk premium), but no longer includes the property tax rate.

Note that the user cost series consists of two terms which are multiplied together: property value, and an expression in parentheses involving interest rates, expected home price appreciation, and tax rates. Movements in the parenthetical expression are dominated by changes in the “gap” between interest rates and expected home price appreciation. We would not expect this parenthetical expression to exceed 20%, nor to drop much below 0%; thus, over long periods of time, user costs must track home prices – the series must be cointegrated. However, “wiggles” in this parenthetical expression will be *strongly* amplified, since the term is multiplied by the home price. Put differently, unless interest rates and home price appreciation move almost perfectly in sync with each other – implying that expected home price appreciation is driven only by interest rates – we should expect user costs to be highly volatile. The larger is γ (or $\tilde{\gamma}$), the less volatile are user costs.

As noted above, I compute three user cost measures. The first is computed using expression (1); it ignores the preferential tax treatment given to homeowners, and assumes that γ equals 5%. The latter two are computed using expression (2), but use different measures of taxes. In one of these latter two, I assume, following Glaeser and Shapiro (2002), that the federal income tax rate is fixed at 0.25, and assume that the property tax is 2%. I assume that the depreciation and maintenance costs are 3%.³ The second user cost measure computed using expression (2) maintains this 2% property tax assumption, but uses the actual marginal federal income tax rate facing a family of four with twice the median income.⁴ The key results are insensitive to these choices, as will be clear below. For the present purpose, I ignore state taxes; these would change the dynamics of the user cost, but would not alter the basic findings of this paper.

To construct these user cost measures, one must forecast four-quarter home price appreciation. The forecasting approach settled upon here (after a fairly intensive search) combines two forecasts: one based upon the previous annual appreciation rate, and the other based upon three lags of quarterly changes in home prices. In each case, only information actually available to agents at time t is used for time t forecasts; coefficient estimates are updated each quarter to reflect the newly-available information. In particular, in each quarter t , two regressions are run. The first regresses the four-quarter appreciation rate on the same variable, lagged four quarters. The time- t forecast generated from this regression uses this coefficient, multiplied by the most recent four-quarter appreciation rate, to produce the expected four-quarter appreciation rate. The second regression has the same left-hand-side variable, but its right-hand-side variables consist of three lags of *quarterly* home price appreciation rates, again lagged four quarters. The time- t forecast generated from this regression applies these coefficients to the three most recent quarterly appreciation rates. The combined forecast takes the simple average of these two forecasts.⁵ Other forecasting approaches were attempted, but none provided more reasonable forecasts.

³ All results are rather insensitive to the choice of these constants.

⁴ These tax rates were constructed, and graciously shared, by Elaine Maag; see Maag (2003).

⁵ Prior to 1979, only the forecast based upon the previous annual appreciation rate is used.

4. Results 1: Graphical evidence

Figure 1 displays the actual home price appreciation, and its forecasted value. The forecast lags the actual home price appreciation (particularly early in the sample), reflecting the surprising amount of persistence of this series. The forecast series, in short, is not too unreasonable. Note that, contrary to what one might expect, the forecast series is *not* appreciably smoother than the actual home price appreciation series. Put differently, appreciation rates – at least at the national level – appear to be extremely persistent; in stark contrast to the behavior of most financial assets, periods of home price appreciation tend to be followed immediately by periods of additional home price appreciation.




Figure 1

Figure 2 displays the mortgage interest rate along with the appreciation forecast. It is clear that these series do *not* move perfectly in sync with each other. This immediately implies that user costs will be highly volatile, as seen below.




Figure 2

Figure 3 presents the estimates of the three different user cost measures. To recap: the first user cost measure ignores deductibility, the second assumes fixed tax rates over the period, and the third assumes that the relevant federal tax rate is the marginal rate facing a family of four with twice the median income. Evidently the tax treatment of mortgage interest payments resulted in negative user costs between 1977 and 1979 for many US homeowners. Since 1981, user costs have fluctuated substantially, but have no upward trend: the steady upward trend in home prices has been effectively “cancelled out” by a reduction (over this period) in the gap between the mortgage interest rate and the expected home price appreciation. As rental prices have risen steadily over this period, this implies that the relative price of homeownership to renting has fallen *substantially* over the period – an issue to which I return below. The decline in the relative price of homeownership is consistent with the concurrent small uptick in homeownership rates.




Figure 3

Figure 4 compares the movement in the two rental series, versus the movement in (simple) user costs. Each series is logged. Then the CPI rent series, and the user cost series, are shifted by a constant so that each has the average value of 1.0 over the sample. The detached rental series is shifted by a constant so that its value in 1988:1 equals that of the CPI rent series.

This graph provides no evidence in favor of the hypothesis that these alternative measures of the cost of housing are measuring the same thing. Rent steadily and smoothly increases over the period. User costs, in contrast, are anything but smooth, and do not appear to share the trend in rents. They jump up dramatically between the late 1970's and 1981, and are more or less trendless after that.⁶ Housing prices also increased steadily over this latter period; the reduction in the differential between mortgage interest rates and expected home price appreciation over this period is responsible for the failure of user costs to track the rise in home prices




Figure 4

⁶ This is true regardless of which of the three user cost series computed here is being considered.

What is particularly startling is the extended period of time during which these two series diverge markedly. To emphasize this point, a 5-quarter moving average of log user cost is plotted in Figure 5 below against the log rent series. The divergent behavior of the two series is even more clearly evident.




Figure 5

The rent inflation series and the user cost inflation series are also quite different. Figure 6 indicates how the volatility of inflation in user costs is vastly larger than that in rents. The difference in volatility is dramatic: the standard deviation of the rent inflation series is 0.00003, compared with that of the user cost series, which is 0.04273. This is a difference of four orders of magnitude.⁷ User costs of homeownership are so volatile that their inclusion in consumer price indices would essentially render such indices useless.




Figure 6

⁷ If the detached rent series is used in the comparison, the difference in volatility is again on the same order of magnitude.

5. Results 2: Regression analysis

The above graphs indicate an empirical puzzle. Rents should be driven by the same underlying forces as user costs: when user costs are high, it is costly to have equity tied up in housing, and rents should rise ... and vice versa: when housing price appreciation is expected to be high, a landlord need not capture as high a rent to be compensated for her opportunity cost of capital. Yet the two series appear to have little relation to each other.

In an attempt to discover whether the two series are as unrelated as they appear, the following regressions were run, on both the aggregate CPI rent series, the CPI rent series corresponding to the four Census regions, and the single-detached rent series:

$$(1) \quad rent_t = \eta + \sum_{j=1}^s \alpha_j rent_{t-j} + \beta i_t + \gamma p_t^H + \phi vac_{t-1} + seasonals + dummies + e_t$$

and

$$(2) \quad \Delta rent_t = \eta + \sum_{j=1}^s \alpha_j \Delta rent_{t-j} + \beta \Delta i_{t-1} + \gamma \Delta p_{t-1}^H + \phi \Delta vac_{t-2} + seasonals + dummies + u_t$$

where $rent$ refers to log rent (and $drent$ its first difference), i refers to the mortgage interest rate (and di its first difference), p^H the log price of housing (and dp^H its first difference), and vac_t the rental vacancy rate (and $dvac_t$ the first difference of the vacancy rate)⁸. Dummy variables for specific periods (such as 1986:II) are included when inspection of residual plots suggests that the residual associated with this time period is an outlier. The second regression imposes a unit root on the rent process – for these persistent series, it is difficult to reject the null hypothesis of a unit root. (Indeed, it is even difficult to reject an I(2) null hypothesis.) Rent lag length s is, in each case, chosen using step-down testing: starting with 5 lags, I reduce the number of lags until the largest lag's estimated coefficient is statistically significant at the 5% level, unless residual correlation suggests that an additional lag is necessary. Variables enter with a lag, rather than contemporaneously, if that specification appears to be preferred by the data.

Table 1 reports the estimated coefficients and their standard errors (in parentheses). Coefficient estimates in bold are significantly different from zero at the 5% level.

⁸ The vacancy rate data is from the Census CPS rental vacancy survey.

Table 1: Rents and their determinants: specification (1)
(|t-statistics| appear below coefficient estimates)

Coefficient	CPI rent index	Detached rent index ¹	CPI rent (NE) ²	CPI rent (MW) ³	CPI rent (South)	CPI rent (West) ⁴
η (constant)	0.029 (5.47)	-0.041 (1.67)	0.027 (2.30)	0.018 (0.015)	0.013 (0.99)	0.019 (1.19)
α_1 (lag of rent)	1.31 (14.32)	0.70 (6.56)	0.84 (9.85)	0.96 (0.08)	0.94 (8.96)	0.94 (11.75)
α_2 (lag of rent)	-0.34 (3.88)	0.63 (4.67)	0.15 (1.84)	0.11 (0.11)	0.28 (2.02)	0.16 (1.40)
α_3 (lag of rent)	--	-0.35 (3.62)	--	0.11 (0.11)	-0.16 (1.15)	0.14 (1.28)
α_4 (lag of rent)	--	--	--	-0.18 (0.07)	-0.30 (2.27)	-0.23 (3.29)
α_5 (lag of rent)	--	--	--	--	0.20 (2.19)	--
β (<i>i</i>)	0.027 (2.11)	0.015 (0.42)	0.063 (2.44)	0.058 (2.16)	0.035 (1.72)	0.045 (1.59)
γ (p^h)	0.025 (4.75)	0.011 (1.88)	0.005 (1.13)	0.006 (0.93)	0.040 (3.72)	0.000 (0.08)
ϕ (vacancy rate)	-0.0010 (2.08)	-0.0012 (1.48)	-0.0021 (3.83)	-0.0025 (4.09)	-0.0011 (3.16)	-0.0026 (4.09)
seasonals	sig.	sig.	sig.	sig.	sig.	sig.
Durbin-Watson statistic	2.06	2.31	1.83	1.81	1.89	2.04
AR(1) of residuals	-0.04 (0.39)	-0.19 (1.55)	0.08 (0.84)	0.09 (0.99)	0.02 (0.25)	-0.02 (0.26)
Jarque-Bera p-value	0.12	0.99	0.53	0.89	0.30	0.47

Notes:

1. Data series runs from 1988:1 – 2003:4.

2. Specification included a dummy for 1986:II. Estimating (1) without this dummy variable, $\hat{\beta}$'s t-statistic drops to 1.56, and the Jarque-Bera p-value drops to 0.00.

3. Specification includes dummy variables for 1979:III, 1979:IV, 1980:IV and 1986:II.

Without these dummy variables, the statistical significance of both $\hat{\beta}$ and $\hat{\gamma}$ change: $\hat{\beta}$'s t-statistic drops to 0.51, while $\hat{\gamma}$'s rises to 2.11. Furthermore, the Jarque-Bera p-value drops to 0.00.

4. Specification includes dummy variables for 1979:III, 1979:IV, and 1980:IV. Estimating (1) without these dummy variables, $\hat{\gamma}$'s t-statistic rises to 2.11, and the Jarque-Bera p-value drops to 0.00.

Table 2: Rents and their determinants: specification (2)
(|t-statistics| appear below coefficient estimates)

Coefficient	CPI rent index	Detached rent index ¹	CPI rent (NE) ²	CPI rent (MW) ³	CPI rent (South) ⁴	CPI rent (West) ⁵
η (constant)	0.001 (1.41)	0.0045 (3.46)	0.003 (2.50)	0.003 (3.49)	0.002 (2.03)	0.002 (1.53)
α_1 (lag of rent)	0.37 (3.84)	-0.15 (1.41)	0.04 (0.41)	0.08 (0.96)	0.12 (1.24)	0.07 (0.78)
α_2 (lag of rent)	0.10 (1.04)	0.40 (3.76)	0.26 (3.11)	0.39 (6.07)	0.46 (5.84)	0.16 (1.97)
α_3 (lag of rent)	0.16 (1.56)	--	0.22 (2.67)	0.36 (4.65)	0.25 (2.97)	0.23 (2.93)
α_4 (lag of rent)	0.28 (2.99)	--	0.24 (2.87)	--	--	0.37 (4.46)
β (i)	0.0338 (0.83)	-0.127 (1.98)	0.184 (2.59)	0.112 (2.64)	0.057 (1.03)	-0.138 (1.73)
γ (p^h)	0.061 (2.32)	-0.0012 (0.03)	0.009 (0.34)	-0.001 (0.02)	0.062 (1.43)	0.113 (2.74)
ϕ (vacancy rate)	-0.0009 (1.15)	-0.0016 (1.97)	-0.0026 (0.53)	0.0001 (0.02)	-0.0044 (0.82)	-0.0133 (2.26)
seasonals	sig.	sig.	sig.	sig.	sig.	sig.
Durbin-Watson statistic	2.03	2.23	1.90	2.02	2.08	1.98
AR(1) of residuals	-0.02 (0.21)	-0.01 (0.07)	0.04 (0.45)	-0.03 (0.30)	-0.06 (0.55)	0.01 (0.09)
Jarque-Bera p-value	0.63	0.64	0.08	0.26	0.41	0.06

Notes:

1. Data series runs from 1988:1 – 2003:4.
2. Specification included a dummy for 1986:II. Estimating (1) without this dummy variable, $\hat{\beta}$'s t-statistic drops to 1.23, and the Jarque-Bera p-value drops to 0.00.
3. Estimated from 1979:02 onwards. Specification includes dummy variables for 1986:II and 1986:III. Without these dummy variables, $\hat{\beta}$'s t-statistic drops to 2.19, and the Jarque-Bera p-value drops to 0.00.
4. Specification included a dummy for 1986:II. Without this dummy variable, the Jarque-Bera p-value drops to 0.00.
5. Specification includes dummy variables for 1986:II and 1986:III. Without these dummy variables, $\hat{\phi}$'s t-statistic drops to 0.73, and the Jarque-Bera p-value drops to 0.00. If one includes these dummy variables, and in addition includes dummy variables for 1979:III, 1980:IV, and 2002:III, $\hat{\gamma}$'s t-statistic drops to 1.50, but the Jarque-Bera p-value rises to 0.55.

Rents are quite persistent, and exhibit significant seasonality. The residual tests performed suggest that the specifications are reasonable. Aside from these noteworthy features of the data, what is most surprising is the “dog that did not bark:” there is only modest evidence that rents respond to the determinants of user costs, and only slightly more evidence that they respond to vacancy rates. (Note that there is no structural model being estimated here; statistically-significant coefficients should be interpreted as indicative of conditional correlations, or as providing conditional improvements in predictive ability.) Thus, it is not surprising that rents and user costs diverge significantly.

1. **Mortgage interest rates** There is evidence that mortgage interest rates significantly influence two of the rental price series investigated: CPI rent in New England, and CPI rent in the Midwest. There is some evidence that the interest rate influences the aggregate CPI rent series (a borderline significant coefficient estimate in specification (1)); however, this does not carry over when the interest rate series is the 1-year US T-bill rate.

There is weaker evidence that the interest rate influences the detached rent index (a borderline insignificant coefficient estimate in specification (2)). Interestingly, the variable enters with the opposite sign. (Furthermore, this result is fragile: if vacancies are either dropped from the specification or enter at another lag, or a trend is included,⁹ this coefficient estimate becomes statistically insignificant.) Collinearity does not appear to be the culprit, since $\hat{\beta}$ is not significantly different from zero even if p^H is dropped from specification (2).

2. **Home prices** There is stronger evidence that increases in home prices increase rents: the aggregate CPI rent index is significantly influenced by changes in house prices. Yet even here there are caveats. First, there is much less evidence that this appears in the more disaggregated series. Second, this effect is evidently much weaker for the detached rent series, where – as with interest rates – we might expect the effect to be stronger.¹⁰

3. **Vacancy rate** This variable enters significantly in all the specification (1) CPI rent index regressions. Surprisingly, it enters significantly in only one of the specification (2) regressions (though there is some evidence of an influence on the detached rent index in that case). Several

⁹ In specifications in which the vacancy rate at lag 2 is not present, inspection of the residuals suggests that a trend should be included.

¹⁰ If the interest rate is dropped from this regression – justified, perhaps, on the basis of its insignificant t-statistic – the t-statistic on $\hat{\gamma}$ rises to 2.11. However, if one drops the vacancy rate as well, $\hat{\gamma}$'s t-statistic falls to 1.33.

authors (e.g., Goodman and Belsky, 1996; Emrath, 2004) note that locating evidence for an influence of vacancy rates on rents is difficult, despite a wealth of anecdotal evidence to the contrary (see, e.g., Max, 2004).

4. Expected appreciation and user costs

The evidence that rents respond to expected price appreciation, or to user costs, is slim. Inclusion of expected aggregate price appreciation in specification (1) for the two aggregate rent series resulted in statistically insignificant coefficient estimates. Including the first difference of this variable in specification (2) resulted in a marginally statistically significant coefficient estimate¹¹ only for the detached rent series.

To examine whether user costs influence rents, specifications (1) and (2) were altered: that interest rates and home prices were dropped as explanatory variables, and user costs were added. There is no evidence that user costs influence the detached rent series, even though – again – this is where we might most expect such an effect to be evident. There is some evidence that aggregate CPI rents respond to user costs – the t-statistic associated with the user-cost coefficient estimate was 2.4 in specification (1) – but this coefficient estimate had a negative sign! (It's t-statistic in specification (2) was a mere 1.4.)

The differential between the mortgage interest rate and expected home price appreciation is, of course, key to understanding user cost movements. A regression of expected home price appreciation on a constant and on the mortgage interest rate suggests that the expected home price appreciation is, on average, $3.5\% + 0.27*(\text{mortgage interest rate})$, or (without a constant) $0.60*(\text{mortgage interest rate})$. Thus, the series are definitely strongly related; however, their correlation is far from 1.

¹¹ t-statistic = 2.13. Interestingly, in this specification the t-statistics associated with both $\hat{\gamma}$ and $\hat{\phi}$ rise above 2.0 as well.

6. Conclusions

1. An ex-ante user cost series, even if smoothed to a great extent, would still be *dramatically* volatile. On pragmatic grounds, this probably means that it should not be used as a measure of homeowner costs in the CPI.¹² Why? Because inclusion of user costs in the CPI would probably render the CPI useless for most purposes.

2. Rent changes are *not* good estimates of homeowner user-cost changes. This means that a major theoretical justification for using rental equivalence in consumer price indices is, on empirical grounds, decisively rejected. Statistical agencies must think about, and be able to justify, the continued use of rental equivalence on other grounds.

For example, statistical agencies might argue that what they are concerned about measuring is simply the rental services, irrespective of how this differs from actual costs of homeownership. Evidently the highly volatile (expected or actual) capital appreciation movements that directly alter homeownership costs do not immediately or directly alter the costs of market rents. There are two related but distinct theoretical arguments that statistical agencies might make in defending the use of a rental-equivalence approach to measuring the costs of homeownership. First, such agencies could argue that, as such appreciation amounts to capital gains or gains to saving, such movements must be excluded as being out-of-scope for a consumer price index. Second, such agencies could argue that such appreciation is instead a component of *income* rather than *costs* – the homeowner as a consumer rents housing from the homeowner as producer, and the capital gains (and some portion of the changes in other parts of the user cost) translate into a change in *production costs*, and hence profits or income, of the homeowner. However, consumer price indices do not attempt to track incomes or profits, merely costs.

¹² Would the user-cost measures that obtain from more realistic models which model adjustment costs feature such dramatic volatility? I suspect that they would. Such user costs will differ from the frictionless user costs outlined in section 3 in that, in place of the expected appreciation term, there will be a term reflecting the *probability of adjustment and realization of the capital gain*. However, the basic driving forces of such user cost measures will still be important: home prices, interest rates, and expected home price appreciation. Home price appreciation is extremely persistent, and mortgage interest rates are not terribly tightly related to this appreciation. To keep a divergence between these variables from translating into a large movement in user costs, the probability of moving would have to adjust in just the right manner. This doesn't seem plausible. For example, consider a period of sluggish interest rates and a sudden increase in home price appreciation, such as occurred during 2003. A standard user cost measure would fall dramatically during an episode like this. Only an equally large *decrease* in the probability of moving would keep a more realistic user cost measure from falling dramatically.

3. Further research is needed to understand what drives rents, and why they diverge from user costs. There is likely some interesting industrial organization work to be done.

4. The BLS is currently investigating the use of user costs for purposes of choosing *weights* to be used in its consumer price index. Given that user costs are so volatile, it is unclear that one would want to use the weights from one particular year to use in the CPI. Thus, the BLS would have to think carefully about what sort of *average* to use; a 10-year forecast average might be the sort of thing settled upon.

References

- Calhoun, Charles. (1996). "OFHEO House Price Indices: HPI Technical Description." Manuscript, OFHEO.
- Case and Shiller (1987) *New England Economic Review*.
- Case and Shiller (1989) "The efficiency of the market for single-family homes." *The American Economic Review* 79, 125-137.
- Diaz, Antonia, and Maria Luengo-Prado. (2004) "On the User Cost and Home Ownership", Manuscript, Northeastern University.
- Dougherty, Ann, and Robert Van Order (1982). "Inflation, Housing Costs, and the CPI." *The American Economic Review*, March.
- Emrath, Paul (2004). Personal communication.
- Gallin, Joshua (2003). "The long-run relationship between house prices and income: evidence from local housing markets." Mimeo, Federal Reserve Board.
- Gillingham, Robert (1980). "Estimating the User Cost of Owner-occupied Housing." *Monthly Labor Review* 103.2, 31-35.
- Gillingham, Robert (1983). "Measuring the Cost of Shelter for Homeowners: Theoretical and Empirical Considerations." *The Review of Economics and Statistics* 65.2, 254-265.
- Glaeser, Edward, and Jesse Shapiro (2002). "The benefits of the home mortgage interest deduction." Mimeo, Hahvahd.
- Goodman, Jack, and Eric Belsky (1996). "Explaining the Vacancy Rate - Rent Paradox of the 1980s." *Journal of Real Estate Research* 11.3, 309-323.
- Maag, Elaine. (2003) "A brief history of marginal tax rates." *The Urban Institute* (www.urban.org).
- Martin, Robert. (2004). Paper on user costs and rents. Manuscript in progress, Board of Governors of the Federal Reserve System.
- Max, Sarah. (2004) "For rent, cheap." *CNN Money*, February 11.