**STICKY RENTS AND THE CPI FOR OWNER-OCCUPIED
HOUSING**

A dissertation
Submitted to
the Temple University Graduate Board

in Partial Fulfillment
of the Requirements for the Degree
DOCTOR OF PHILOSOPHY

by
Adam Ozimek
August 2013

Examining Committee Members:

Dr. Moritz Ritter, Advisory Chair, Department of Economics
Dr. William Dunkleberg, Department of Economics
Dr. Forrest Huffman, Department of Finance
Dr. Richard Voith, External Member, Econsult Solutions Inc.

©
Copyright
2013

by

Adam Ozimek
All Rights Reserved

ABSTRACT

This dissertation examines the implications of sticky rents on the measurement of owner-occupied housing in the Consumer Price Index (CPI). I argue that marginal and not average rents are the most theoretically justified measurement of owners' equivalent rent (OER), and that the current measurement of rental inflation using average rents is methodologically incorrect. I then discuss the literature on sticky rents and tenure discounts and present a theoretical model showing the implications of sticky rents for aggregate measures of inflation. Then I use two new data sources to construct marginal rent measures to compare to average rent measures. The results show that marginal rents reflect market turning points sooner, and show a larger post-housing bubble decline in rents. In addition, marginal rents are shown to forecast overall inflation better than average rents. Finally, the implications of these results for policy are considered using the Taylor Rule for optimal monetary policy. The results present suggestive evidence that the impacts of switching to marginal rents may be large enough to significantly impact monetary policy and allow the Federal Reserve to be more responsive to both the boom and bust of housing bubbles.

ACKNOWLEDGEMENTS

I have benefitted tremendously in this research from the advising of my committee chair Moritz Ritter, without whom this would have been a much, much lower quality dissertation. Committee members William Dunkleberg and Forrest Huffman also provided much useful advice and comments. Special thanks goes to Richard Voith for assistance both as my external reader, my employer, and as the first person to raise the issue to me of whether marginal or average rents were what the CPI for owner-occupied housing should be measuring. Likewise I have been helped greatly by conversations and other assistance on this issue from Peter Angelides, Randal Verbrugge, Paul Sullivan, Peter Linneman, Steve Mullin, David Crawford, David Genesove, Brian Glassman, Dimitrios Diamantaras, Charles Swanson, James Bailey, Jerry Hionis, Claudia Kurz, Jay Lybik, Andrew Strauch, Larry Wisdom, Craig Thomas, Svenja Gudell, Stan Humphries, and Jed Kolko. Last but not least, the following would be nearly unreadable were it not for the patient editing of my wife Amelia Ozimek.

TABLE OF CONTENTS

ABSTRACT	iii
ACKNOWLEDGEMENTS	iv
LIST OF TABLES	vi
LIST OF FIGURES	vii
CHAPTER	
1 INTRODUCTION	1
2 CURRENT BLS MEASUREMENT	5
2.1 OER Expenditure Weight	5
2.2 Survey and Sampling Procedure	7
2.3 Adjustments	9
2.4 Estimating the Price Relative	11
3 MARGINAL VERSUS AVERAGE RENT AND OER	13
3.1 The Theoretical Justification for OER	13
3.2 User-Cost	17
3.3 Weaknesses of User Cost as a Measure of Housing Services	19
3.4 User Cost as a Rental Equivalence Justification	20
3.5 Does User Cost Justify Marginal or Average Rents?	23
3.6 Opportunity Cost	25
3.7 Does Opportunity Cost Justify Marginal or Average Rents?	29
4 STICKY RENTS AND TENURE DISCOUNTS	32
4.1 Sticky Rents	32
4.2 Tenure Discounts	36
5 IMPLICATIONS OF CURRENT OER MEASUREMENT	40
5.1 Effect of Nominal Rigidity on Rent Levels	40
5.2 Effect of Nominal Rigidity on Rent Changes	42
5.3 Effect of Nominal Rigidity on the CPI Measure of Inflation	44
5.4 Changes in Turnover	47
5.5 Endogenous Price Changes	50

5.6	Additional Lags in Current CPI Measurement	55
5.7	Summary	56
6	ALTERNATIVE MEASURE OF OER	57
6.1	Measuring Marginal Rents	57
6.2	Lease Level Rental Data	60
6.2.1	Data Description	60
6.2.2	Inflation Comparison	63
6.3	MLS Rental Data	67
6.3.1	Data Description	67
6.3.2	Inflation Comparison	72
6.4	Repeat-Rent Shortcomings and Extensions	76
6.5	Policy Implications	80
6.5.1	Overview	80
6.5.2	Forecasting Comparison	81
6.5.3	Taylor Rule	86
6.6	Conclusions	93
	BIBLIOGRAPHY	95

LIST OF TABLES

2.1	CPI Expenditure Groups and Weights	6
4.1	Average Lease Length	33
4.2	Probability of Rent Change	34
4.3	Probability of Twelve-Month Rent Change by Tenure	37
6.1	Unit Level Characteristics	61
6.2	Repeat-Rent Observations in PSU A312	71
6.3	Count of Transactions by Number of Repeat-Rents	72
6.4	Effect of Removing Single Sales on Sample Size	78
6.5	Measures of Forecasting Performance	86

LIST OF FIGURES

2.1	87 Primary Sampling Units	7
3.1	Rent Measure Comparison	30
4.1	Level of Rents by Tenancy Duration	38
5.5	Simulations of Tokyo CPI Under Alternative Assumptions	54
6.1	Histogram: Duration of Price Spells	63
6.2	Hazard Function: Percent of Units With Zero Change	64
6.3	Rent Index Method Comparison	66
6.4	Rent Index Method Comparison, 12-Month Percent Change	68
6.5	Repeat-Rent Index Compared to CPI for OER	73
6.6	Repeat-Rent Compared to CPI for OER One-Month Changes	74
6.7	Smoothed Repeat-Rent Compared to CPI for OER One-Month Changes	75
6.8	Repeat-Rent Compared to CPI for OER Twelve-Month Changes	76
6.9	Repeat-Rent Compared to CPI for OER and Case-Shiller for D.C., Year-on-Year Changes	77
6.10	Basic and GLS Repeat-Rent Comparison Year-on-Year Percent Changes	79
6.11	Forecast Comparison $\pi_{t,t+h}$ = OER, Washington-Baltimore CMSA	83
6.12	Forecast Comparison $\pi_{t,t+h}$ = Overall CPI, National	84
6.13	Forecast Comparison $\pi_{t,t+h}$ = Overall CPI, South Region	84
6.14	Overall CPI Comparison 12-Month Average Changes, D.C./Baltimore CMSA	88
6.15	Percentage Point Change in Taylor Rule Implied Federal Funds Rate From Adopting Marginal Rent, Using MLS Data	90
6.16	Percentage Point Change in Taylor Rule Implied Federal Funds Rate From Adopting Marginal Rent, Using REIT Data	92

CHAPTER 1

INTRODUCTION

Shelter makes up 32% of the consumer price index (CPI), with the owner-occupied housing (OOH) portion comprising 24% and rent of primary residence making up 6%¹. This means housing is the single largest component in the CPI. Reflecting this importance, the current Bureau of Labor Statistics (BLS) methodology has come under some criticism as a large U.S. house price bubble and crash occurred alongside slow and steady growth in the CPI for OOH. Part of the reason for this divergence may be explained by the current BLS approach to measuring owner-occupied housing inflation.

This dissertation will argue that, based on the theoretical foundation of price indexes, the CPI for OOH should use current market rents instead of the average rent across all surveyed households that the BLS currently uses. I will provide theoretical and empirical evidence that, due to the nominal rigidity of housing rents, the mismeasurement resulting from the current approach is consequential. In addition, I propose an alternative and conceptually preferable measurement method and investigate effects of choosing this alternative method over the status quo.

This study begins by providing an overview of the current BLS method of using rental equivalence to estimate the CPI for OOH, beginning the sampling methodology, continuing with various rent adjustments that are made, and finally to the eventual estimation of OOH inflation. This section illustrates that the BLS is currently using average rents to measure OOH inflation.

Having established what the CPI is currently measuring, Chapter 3 addresses whether the BLS measurement is in line with their stated goals. To do this, Chapter 3 introduces the concept of marginal rents as a potential alternative to the current

¹Percentages are for the the CPI-U as of December 2010. A remaining 0.78% of shelter is lodging away from home.

measurement. Marginal rents reflect current market conditions and are based on prices that a marginal buyer in the market would face. This is in contrast to average rents, which are based on average household rent expenditures. The main cause of divergence between these two measures is that rents are typically set by long-term leases. This means that the current market price that a unit would rent for, which would be used in marginal rents, may not be the same as the current rent paid on that unit, which would be used in average rents. Average rents are often set months or even years in the past while marginal rents reflect current prices.

The main contribution of Chapter 3 is to present a thorough argument that marginal rents are more consistent with both of the primary theoretical justifications of rental imputation: the user cost justification and the opportunity cost justification. While some discussion of marginal versus average rents exists in the literature, this is the first study to provide a rigorous case for marginal rents grounded in the theoretical justification for rental equivalence.

The conclusions of Chapter 3 suggest that the current method is mismeasuring rents. Chapter 4 then broadens the discussion of the underlying reasons why the choice of marginal versus average rents may be significant: the micro-rigidity of rents. This chapter summarizes the existing empirical evidence that rents adjust infrequently, i.e. that they are "sticky," and that tenant tenure discounts mean existing tenants pay lower rents than new tenants. Therefore, average rents may not be equal to marginal rents.

While the paper thus far has suggested that marginal rents may diverge from average rents due to the micro-rigidity discussed in Chapter 4, there has been no evidence that micro-rigidity will translate to macro-rigidity, and therefore aggregate measures of average rents may still equal marginal rents. In Chapter 5, I present theoretical evidence that long and short-run divergences can occur between aggregate measures of average and marginal rents, and the resulting CPI inflation. Motivated by infrequent

rent renegotiation and long-term contracts observed in the data, a Calvo model of price setting is used to show the relationship between aggregate average and marginal rents in levels, changes, and inflation rates. In addition, illustrative simulations show how one-time shocks can affect marginal and average rents differently. Importantly, these results indicate that marginal rents will reflect underlying shocks in a more timely manner than average rents.

This section concludes with a discussion of short-comings of the model and a review of the literature on the relationship between micro-rigidity and macro-rigidity for housing. The available evidence suggests that empirical estimates are needed to assess the extent that micro-rigidity impacts macro-rigidity. This issue is addressed in Chapter 6 which contains empirical estimates of the implications of marginal versus average rents. First, a method for measuring marginal rents is proposed. Then a new, proprietary dataset comprised of 397,137 leases for 81,842 multifamily units over the past decade from a large Real Estate Investment Trust (REIT) is introduced. This data is used to estimate rental inflation using both a proxy of the existing CPI methodology of average rents and the proposed marginal rent methodology. The results provide empirical evidence of the extent to which marginal and average rents diverge. Overall, the marginal rent methodology shows inflation and deflation significantly sooner and more starkly than the CPI method. In particular, the deacceleration in prices coinciding with the bursting of the house price bubble shows up in year-on-year changes twenty-six months earlier in the marginal rent than in the official CPI measure. In addition, deflation appears in the marginal rent series nine months earlier than in the official CPI measure.

As an additional test of the significance and preferability of marginal rents, Chapter 6 introduces a dataset of 276,158 single family housing rental listings from 2000 to 2012. This data comes from a regional Multiple Listing Service (MLS), which are databases used by Realtors to find and list properties, and are a common source

of home sales data in real estate economics. In contrast to the multi-family REIT data, this dataset has detailed geographic information, a large percentage of detached housing, and contains a significant overlap with a particular BLS sampling area. This allows the estimation of a marginal rent inflation series that can be compared to the actual BLS CPI for OOH in this geography. The marginal rent series shows the popping and bottom of the housing bubble more quickly than the CPI series, however it does not show an overall greater rise in prices during the bubble period. The deacceleration of year-on-year price changes shows up two months earlier in the marginal rent series. The bottom of the bubble shows up even sooner in the marginal rent series, with the largest decline in prices showing up fourteen months earlier than in the CPI. Overall, the marginal rent series better reflects the Case-Shiller House Price Index for the region, but OER still lags far behind house prices suggesting measurement methodology does not account for the majority of the divergence between house prices and rents. Following the literature on inflation measurement, the forecasting ability of marginal rent and current OER measures are compared. The evidence suggests marginal rent is better able to forecast overall inflation, however this may largely be due to the ability of marginal rent to forecast the current BLS measure of OER better than the BLS measure forecasts itself.

The CPI is one of the most important measures of inflation and is followed closely by the Federal Reserve, businesses, and governments. Reflecting this, potential biases and measurement problems with the CPI have long been a topic of research interest. This study's contribution is to extend this literature. In addition, because the evidence suggests that the mismeasurement discussed in this paper causes real underlying OOH price changes to show up slowly in measured inflation, then correcting this may allow the Federal Reserve to be more responsive to house price bubbles and bursts, thereby responding to critics' concerns without altering its existing goals and targets.

CHAPTER 2

CURRENT BLS MEASUREMENT

The Bureau of Labor Statistics measures inflation in owner-occupied housing differently than any other component in the Consumer Price Index. Instead of measuring prices from current housing sale prices, the BLS imputes the price of OOH using rents. This is a departure from all other components of the CPI which are measured using current market prices of goods and services. The justification for this imputation is an important consideration in addressing measurement issues and will be reviewed in a later section. First, the following section will provide a detailed discussion of the current measurement procedure.

2.1 OER Expenditure Weight

The highest level decision about owners' equivalent rent (OER) and all components of the CPI is what weight they should be given in the overall CPI. The CPI categorizes household spending into 70 expenditure categories within 8 major groups, and at the most refined level there are 211 item strata. Housing is one of the 8 major groups within which Shelter is an expenditure category within which OER of primary resident is an item strata. As the following table shows, this strata receives nearly 23% of the the CPI weight, making it the single largest component.

The weights for each category are derived from the Consumer Expenditure Survey, which is a quarterly survey of households performed by the U.S. Census Bureau. For housing, the following question is asked of renters:

What is the rental charge to your Consumer Unit for this unit including any extra charges for garage & parking facilities? Do not include direct payments by local, state or federal agencies. What period of time does this cover?

Table 2.1: CPI Expenditure Groups and Weights

Expenditure Group	Weight
Food and beverages	15.256
Housing	41.02
-Shelter	31.539
–Rent of primary residence	6.485
–Lodging away from home	0.749
–Owners’ equivalent rent of resident	23.957
— Owners’ equivalent rent of primary residence	22.543
–Tenants’ and household insurance	0.348
-Fuels and utilities	5.372
-Household furnishings and operations	4.109
Apparel	3.562
Transportation	16.875
Medical care	7.061
Recreation	6.044
Education and communication	6.797
Other goods and services	3.385
Total	100

And the following is asked to homeowners:

If someone were to rent your home today, how much do you think it would rent for monthly, unfurnished and without utilities?

These questions are used only to determine weights for the housing categories (Poole, Ptacek, Verbrugge, 2005). Previously, these weights were only updated every 10 years. However, beginning in 2002 the BLS moved to a 2 year updating period, so the index was a closer approximation to an ideal cost-of-living index rather than a Laspeyres index (Greenlees and Williams, 2009). For example, the weights as of December 2011 were based on the 2008-2009 CES (BLS, 2012).

2.2 Survey and Sampling Procedure

The data source for residential rents is the CPI housing unit sample, a survey conducted specifically for this purpose.¹ Like other CPI components, the largest sampling geography is the 87 primary sampling units (PSUs) shown in Figure 2.1 below. These PSUs are used to generate the price indices for 38 CPI Index Areas. Thirty-one of the PSUs are “self-representing,” meaning an individual PSU represents an individual Index Area. The other 56 PSUs are “non-self-representing”, which collectively represent the remaining 7 index areas. For example, one Index Area represents 38 small Northeast metropolitan areas, including Buffalo, Hartford, Syracuse, Burlington, and others. Eight of these metros were randomly selected to represent all of them, and each of these eight areas is a non-self-representing PSU (BLS Handbook of Methods).

Figure 2.1: 87 Primary Sampling Units

Every PSU is divided into six strata that each represent approximately 1/6 of the total PSU housing expenditure. Within strata, neighborhoods called “segments” are designated which are composed of groups of adjacent census blocks. Each segment must contain 50 housing units for larger PSUs or 30 housing units for smaller PSUs

¹This section will draw from the extensive overview of CPI for housing methodology in Poole, Ptacek, and Verbrugge (2005), and the BLS Handbook of Methods (2012)

and contain on average 150 housing units. Each segment is placed in one of the six strata.

The BLS lists five goals it attempts to accomplish via this geographical stratification. First, geography accounts for most of the variance in rent change, and rent level is the second most important factor. Thus geographic stratification allows sample coverage for these two most important characteristics. Second, it allows for segments with renter units and owner units, which is the basis for the measurement of OER. Third, the design allows new construction be added to existing geographic strata. Fourth, by accomplishing the previous three goals, it should also limit the sampling variance of rent and OER indexes. Finally, stratification structure allows sample rotation on a rolling basis.

Within each strata, a sample of segments is chosen with the odds of being chosen proportional to total housing expenditures in that segment.² Each segment is then assigned to panels that determine when in the year the houses they contain will be interviewed. Each panel is interviewed every six months, and a different panel is interviewed every month. Panel 1, for example, is surveyed in January and July each year while panel 2 is priced in February and August and so forth. The segments are selected into the panels so that each panel, and therefore each month's measurement, is representative.

Each segment is intended to produce five housing units for the sample. For the 1998 housing sample, 50,000 units were desired, so 10,000 segments were chosen. However, due to the low number of renters in some areas, the initial sample was approximately 25,000. To increase the sample size, an augmentation was done to increase the number of houses sampled in segments with 3 or fewer houses. This yielded an additional 10,000 houses. In 2010, the number of housing unit survey responses used in the estimation of the CPI was 57,015, which implies a sample of

²Housing expenditures used are total rents for rental units, and total owners' equivalent rent for OOH.

over 28,000 units (Crawford, Mauro, and Church, 2011).

2.3 Adjustments

Once the contract rents are collected from the units in the sample, several adjustments must be made. There are two measures of rent that must be estimated from the contract rents: “economic rent” for the rental index, and “pure rent” for the owners’ equivalent index. Economic rent is the full amount paid to the landlord, including any subsidies paid to the landlord through government programs like Section 8, and any rent reductions provided in exchange for services the tenant performed for the landlord, with an adjustment for any change in the quality of the housing services. The pure rent for OER is identical to the economic rent less the value of landlord-provided utilities. The subtraction is made because utilities are out-of-pocket expenses for homeowners that are measured separately in the Fuels and Utilities CPI category, and so including them in OER would amount to doublecounting.

Rents for vacant units are imputed using average rent changes within the PSU. For units that have been vacant less than six months, rent is assumed to have increased at the average rate of change for non-vacant units that have a new tenant within the last six months. For units that have been vacant for more than six months, the average rent change for all non-vacant units is used.

There are three adjustments made to control for quality changes. These are in addition to the longitudinal nature of the survey which already controls for a great deal of quality change by looking at the same units over time. The additional adjustments are needed as quality changes for a given unit over time can have a non-trivial effect on measured rents that reflect an improvement in services, not just paying more for the same unit. To account for this, the BLS makes three types of quality adjustments: aging, structural change, and utility changes.

The BLS has performed the adjustment for the aging of the housing stock since

1988, based on the observation that as buildings age they generally tend to decline in value. For example, if the rent on a house goes up 3% from one period to the next, while aging decreases its quality by 1% over the same period, then total inflation should be 4%. Some housing will receive maintenance to counteract the aging based depreciation, but over time most housing units are either eventually torn down or completely rehabilitated, so the housing stock overall must be depreciating (Lane et al, 1988). In order to control for the aging bias effect, the BLS uses hedonic regression models. These models are run annually using every unit in the CPI Housing Survey from the previous year. They model log of rent as a function of:

- 10 structural characteristics: number of bedrooms, bedrooms squared, number of other rooms, other rooms squared, bathrooms, whether the unit is detached, and dummy variables for whether the unit has central air, oil heat, electric heat, and window air conditioners.
- 10 neighborhood characteristic variables, each as a percent of the neighborhood: white, large buildings, two or more cars, air conditioned, children ages 8 to 18, below poverty level, aged 65 and over, mobile homes, some college, and unemployed.
- 3 utilities and services provided in rent: gas, electric, and parking
- 6 depreciation variables: age, age squared, age interacted with a dummy if the built before 1900, age interacted with detached building dummy, and age interacted with number of rooms (Poole, Ptacek, Verbrugge, 2005)

The aging regression is run at the Census Region level (West, Midwest, Northeast, South), and depreciation rates are calculated for each PSU using the marginal impacts of the aging coefficients from the above regression and the PSU average of the relevant variables. Therefore, every unit in a given PSU has the same depreciation rate. This

is done because many units in the housing survey do not have information on year built, which is used in the aging regression.

The second type of quality adjustment is for structural changes. Prior to February 1989, any structural changes lead to a rental unit being left out of the CPI estimate for that month with its weight being redistributed to nearby units. However, because rents are likely to change at the same time quality improvements are made, this potentially biased inflation downward (Henderson and Berenson, 1990). An added benefit of the aging bias regressions is that estimated coefficients for various physical housing characteristics also allow quality adjustments to be made for major structural changes. These adjustments are made for four types of structural changes: central air, the number of bedrooms, bathrooms, and other rooms. The coefficients for these adjustments are updated once a year when the aging bias regression is re-run using the previous year's housing data³.

2.4 Estimating the Price Relative

Each month the BLS estimates a price relative that is used to move the previous month's consumer price index forward. For OER, there are two parts that determine the price relatives for a given area: the weights assigned to each segment and the rents for those segments. For each segment, weighting is based on the aggregate housing rents and aggregate owners' implicit rents from the 1990 Census. The owners' implicit rents are estimated using owner reported housing values as described in the previous section.

The price relatives (PR) for each PSU are estimated as follows:

$$PR_{p,t,t-6} = \frac{\sum_{i \in p} W_s \cdot R_{i,s,t}}{\sum_{i \in p} W_s \cdot R_{i,s,t-6}} \quad (2.1)$$

³The frequency of the adjustment is not stated in the BLS Handbook or other sources. This fact comes from correspondance with Randall Verbrugge of the BLS.

Where $PR_{p,t,t-6}$ is the price relative for period t to $t-6$ in area p , and W_s is the weight for houses in segment s .

The index $I_{p,t}$ is estimated by moving the index $I_{p,t-1}$ forward using the sixth root of $PR_{p,t,t-6}$, which approximates the one month change:

$$I_{p,t} = I_{p,t-1} \cdot \sqrt[6]{PR_{p,t,t-6}} \quad (2.2)$$

CHAPTER 3

MARGINAL VERSUS AVERAGE RENT AND OER

There are two distinct ways to measure prices: marginal prices and average prices. Marginal rents are the prices that prevail in the current market at current prices. Average rents are the prices paid on average in a given period. Marginal prices are what the BLS uses for all CPI items except housing. To be discussed later, there are important reasons why marginal and average rents diverge, but a sufficient condition is the presence of rents set by long-term contracts. Average rents would be estimated, for instance, by taking a survey of household rents, which would be at various stages in long-term contracts. These would not reflect marginal rents because the rents that some households pay in a current period are set by contracts set in the past. Only a subsection of the households will have signed market contracts in the current period and therefore be paying marginal rent prices. In order to measure marginal rents, one would need to survey the current prices for rental units available in the market.

Given the divergence between marginal and average rents, the question arises as to which the CPI should measure for OER. The following sections will discuss the theoretical justifications for OER and argue that, based on cost of living index theory, marginal rents are more appropriate than average rents. This requires a consideration of the two main justifications for rental equivalence: user cost, and opportunity cost. These two theories will be discussed in turn, including the cost-of-living theory that underlies both, why the latter is regarded as the primary justification, and what each implies for the proper measurement of OER.

3.1 The Theoretical Justification for OER

It is not immediately obvious that rents should be used to measure inflation in owner-occupied housing services. Therefore, it is not immediately obvious whether

marginal or average rents should be used. To understand which kind of measurement the BLS should be pursuing, it is necessary to first establish a justification for rental equivalence. In other words, how rents should be measured cannot be answered without first discussing why rents are measured in the first place.

Ultimately, all CPI measurement issues must be considered in light of the stated measurement goal of the CPI, which has explicitly been a cost-of-living index (COLI) since it was recommended by the Boskin Commission in 1995 and shortly after accepted by the BLS (Greenlees, 2006). Even before the BLS explicitly accepted cost of living theory as a measurement goal, it was considered a guide in dealing with operational problems (Greenlees, 2006). The BLS position here is consistent with a broad range of literature that agrees with this point including the Boskin Commission (Gillingham, 1983; Schultze and Mackie, 2002; Boskin Commission, 1996). The intended use of the CPI as a COLI is therefore important driver of the choice of measurement.

A COLI attempts to measure the changes in the cost of achieving a particular level of satisfaction for a given consumer. The relevant satisfaction is assumed by cost of living theory to be a function of the consumer's utility function (Gillingham, 1983). Specifically, the BLS Handbook of Methods defines a COLI for the current month as answering the following question:

What is the cost, at this month's market prices, of achieving the standard of living actually attained in the base period?

The cost of living measurement goal explains when and why inflation for durable goods should be measured differently than for non-durable goods. Usually consumer price indices, including the BLS CPI, only consider expenditures on goods as all occurring in the period that the good is purchased. This is known as the acquisitions approach. The problem with an acquisitions approach for durable goods is that they represents a stock of services, whereas a cost of living theory assumes that welfare is

determined by the flow of services that a consumer receives from the durable good (Gillingham, 1983). Therefore, it is the price of that flow of services, and not the price of the stock, that should be measured in a COLI.

This conclusion implies that all durable goods in the CPI should be measured by the cost of their flow of services, not the price to acquire the asset. However, housing is the only durable good in the CPI not measured using the acquisition approach. Deiwert (2003) points out that this may simply be due to tradition, with price statisticians following the approach taken by Alfred Marshall and others over the past century. Another more defensible explanation Deiwert offers is that the faster the good depreciates, i.e. the shorter its useful life, the more closely the acquisition approach will approximate value of the flow of consumption services. Or, as Gillingham and Lane (1982) phrase it, “the flow of aggregate services is closely related to the flow of aggregate purchases, and service price movements are closely related to asset price movements.” Clothing, for instance, is technically a durable good since the consumer may derive a flow of services from it for more than a year.¹ However, the relatively fast depreciation of clothing means it will likely not yield a useful life of more than a few years on average, and so inflation measured using the acquisition approach will not greatly differ from an ideal user cost approach. The separate treatment of housing is due to the fact that its useful life averages tens of years, and so the depreciation is much slower, and the divergence between acquisition and ideal user cost is potentially much larger.²

¹The System of National Accounts defines a durable as “a good that may be used for purposes of consumption repeatedly or continuously over a period of a year or more”, which Deiwert (2003) interprets more broadly as “it can deliver useful services to a consumer through repeated use over an extended period of time”. The Stigler Commission acknowledges that since most goods are not consumed at the moment of purchase, durable is “an elusive concept” (Price Statistics Review Committee, 1961). They focus on commodities whose useful life is long enough that there is a relatively healthy used market.

²Some, including the Boskin Commission have argued that longer lived durables like automobiles should be considered as having significant enough useful life to justify departing from the acquisition approach, either by adopting a user cost or rental equivalence. Whether or not this is true is not clear, but is beyond the scope of this paper.

In recognition of this fact, the BLS focuses on measuring the cost of consuming housing services. However, the market price of owner-occupied housing services is not directly observable, therefore indirect measurement techniques must be used. There are two primary ways to measure housing services that are considered by price statisticians, and the BLS in particular: user cost and owners' equivalent rent.

An example from Dale Smith (1975) illustrates these two measurement approaches. Consider measuring the cost of food for a family that grows everything in a garden. One way to estimate the cost would be to add together all of the input costs including soil, seeds, and labor. This would be the user cost. A rigorous version of this would include the cost of buying the land and other durable components at the beginning of the period less what one gets from selling it at the end of the period. Another approach would be to say that the cost of consuming the food the family grew is equal to the price they could have sold it for. This would be the rental equivalence approach. Relevant to this paper is the fact that in equilibrium the two measurements should be equal. Therefore, regardless of what one thinks of the opportunity cost justification of rental equivalence (e.g. the idea that the cost of the food is how much it could be sold for), if the user cost is the correct conceptual framework, then the equality of the two justifies using the rental equivalence approach.

As Gillingham (1980) argues, these two approaches have “substantially different operational implications”, but are “conceptually equivalent”. As a measurement method, the user cost approach has fallen out of favor due to a number of issues, including some theoretical issues, but more importantly due to the complexity of actually measuring the components of user cost (Poole, Ptacek, Verbrugge, 2005). However, despite its declining importance as an alternative measure of housing services, the conceptual equivalence between user cost and rental equivalence remains important as it provides one of the two main justifications for rental equivalence. The other justification is “opportunity cost justification” (Poole, Ptacek, Verbrugge,

2005). This takes several conceptual forms, but the underlying concept is that the rental price of a house represents the opportunity cost of owning it (i.e. owners could rent the house out and forego this rent to live there).

The following sections will review the relative strengths and weaknesses of user cost and opportunity cost, as both measurement approaches and as justifications for rental equivalence.

3.2 User-Cost

The most simple formulation of the user cost measurement approach interprets the cost of housing services to be the cost of purchasing a house in the beginning of a period, using it over that period, and then selling it at the end. This “user cost of capital” theory is used in a variety of contexts including capital asset pricing, production function studies, measurement of total factor productivity, and analysis of depreciation (Deiwert and Nakamura, 2009). User cost is derived from the fundamental equation of capital theory, which dates to at least the mid 1800s. This equation states that “in equilibrium, the price of an asset will equal the present discounted value of the future net income that is expected to be derived from owning it”. From this fundamental equation of capital theory, user cost theory derives an implicit market rental price for using an asset in terms of the cost of ownership. (Katz, 2009).

The following simple example, from Katz (2009) illustrates how user cost can be derived from the fundamental equation of capital theory.

First, the purchase price of a house V_t at the beginning of period t is expressed as the discounted present value of the net benefits, which is the rent less the operating costs. Let u_t represent the expected end of period t value of housing services, which is the benefit derived from owning a house in period t . Let o_t be the operating costs of owning the housing unit in period t , to be paid at the end of the period. Note that period t benefits and costs flow at the end of the period, and so the current price will

reflect the value of these goods discounted 1 period. The house has a useful life of n periods. Let the interest rate at time t be r_t . Then the fundamental theorem of capital tells us that in equilibrium:

$$V_t = \frac{u_t}{(1+r_t)} + \frac{u_{t+1}}{(1+r_t)(1+r_{t+1})} + \dots + \frac{u_{t+n}}{\prod_{j=t}^{t+n}(1+r_j)} - \frac{o_t}{(1+r_t)} - \frac{o_{t+1}}{(1+r_t)(1+r_{t+1})} - \dots - \frac{o_{t+n}}{\prod_{j=t}^{t+n}(1+r_j)} \quad (3.1)$$

Shifting ahead one period, the services u_t and costs o_t will be accrued at the end of t , so the house price in period $t+1$ will equal:

$$V_{t+1} = \frac{u_{t+1}}{(1+r_{t+1})} + \frac{u_{t+2}}{(1+r_{t+1})(1+r_{t+2})} + \dots + \frac{u_{t+n}}{\prod_{j=t+1}^{t+n}(1+r_j)} - \frac{o_{t+1}}{(1+r_{t+1})} - \frac{o_{t+2}}{(1+r_{t+1})(1+r_{t+2})} - \dots - \frac{o_{t+n}}{\prod_{j=t+1}^{t+n}(1+r_j)} \quad (3.2)$$

Then dividing the right-hand side of 3.2 by $(1+r_t)$, subtracting it from 3.1, and solving for u_t yields:

$$V_t - \frac{V_{t+1}}{1+r_t} = \frac{u_t}{(1+r_t)} - \frac{o_t}{(1+r_t)} \quad (3.3)$$

$$u_t = r_t V_t + o_t - (V_{t+1} - V_t) \quad (3.4)$$

This is the end of period t user cost. User cost theory thus provides a conceptual basis with which to think about the cost of owner occupied housing, e.g. buying in one period and selling in the next, and also provides a way that the cost of owner-occupied housing can be built up from constituent costs: $r_t V_t$ is what could have been earned had the money on the house been invested elsewhere at the nominal interest rate, e.g. the opportunity cost; o_t is the cost of maintaining the house in that period; and finally $(V_{t+1} - V_t)$ is nominal losses or gains from selling the house, e.g. the price the house is sold for V_{t+1} , less the price paid for it, V_t . This provides price statisticians

a method to, in theory, construct owner-occupied inflation by first estimating these component costs. In fact, the methods used by the BLS from the early 1950s until 1983 represented a form of user cost known as the “asset price approach” which is a subset of the “payments” or “cash flow” approaches that in general attempt to build up a homeownership cost out of average out of pocket expenses (Poole, Ptacek, Verbrugge, 2005).

3.3 Weaknesses of User Cost as a Measure of Housing Services

While user cost theory presents an argument for how the value of the flow of housing services could be measured for the purposes of inflation, user cost measures suffer from a variety of problems in theory and in practice.

One fundamental problem is that user cost may not be consistent with cost of living theory that is the measurement goal of the CPI. Housing is both an asset and a flow of consumption services, and user cost explicitly includes the part of housing that is related to its nature as an asset. Verbrugge (2008) argues housing can be thought of as a bundle of two things: a depreciating durable consumption good, which is the house itself, and an appreciating financial asset, which is the land. The appreciation, interest rates, and some other parts of user cost are characteristics of housing as an asset, rather than housing as a flow of consumption services. Importantly, assets are not considered within the scope of cost of living indexes. As Poole, Ptacek, and Verbrugge (2005) argue:

...it is difficult to justify why the investment returns on one category of assets – namely, the housing unit that the household occupies – should be reflected in the CPI, while other investment returns are excluded.

Paradoxically, these aspects which are outside of the scope of the CPI determine the equilibrium cost of the flow of housing services which is within its scope.

In addition to the theoretical issues with user cost, there are a variety of practical

issues that would complicate the construction of user cost indices. First, empirical estimates of user cost from Verbrugge (2008) showed high volatility with a standard deviation of 0.04273 compared to 0.00003 for a rental equivalence measure. While this does not indicate which measure is more conceptually accurate, Verbrugge argues that the volatility of user cost is such that “their inclusion in consumer price indices would essentially render such indices useless”. With owner-occupied housing inflation constituting approximately a quarter of the CPI weight, a large amount of volatility like this would “drive the entire index on a month-to-month basis, likely drowning out the signal in noise” (Garner and Verbrugge, 2009).

Another problem, which is theoretical and empirical, is that user cost measures are built up from variables which could be specified a variety of different ways, and the best way is not always obvious. In particular, one must determine the appropriate expected appreciation, which means generating forecasted appreciation. This is especially difficult given that there is no agreed upon model of house price dynamics (Verbrugge and Garner, 2009). Specification of the proper interest rate is another challenge (Garner and Verbrugge, 2009). Which interest rate is theoretically desirable is unclear, and the interest rate chosen has substantial impact on the resulting estimates (Gillingham, 1983).

Given the variety of theoretical and empirical problems with user cost measures, researchers have been concluding for at least 30 years that user cost estimates are not a practical solution for the CPI.

3.4 User Cost as a Rental Equivalence Justification

While user cost may not be desirable to estimate the cost of housing services, it has long provided a theoretical justification for an alternative measure: owners’ equivalent rent. From basic Jorgensian capital theory, it can be shown that, in equilibrium, this simple, frictionless user cost will equal the rental price of a housing unit (Van

Order, 1982; Gillingham, 1980; Gillingham, 1983). In fact, early research arguing for OER did so partly based on its supposed equalivalence with user costs (Gillingham, 1983; Gillingham, 1980). Market rents, in this framework, have been viewed as “an appropriate measurement tool for user costs” (Verbrugge, 2004).

However, while a simple model shows that rents and user costs should be equal in equilibrium, a slightly more complex user cost model from Vebrugge (2008) and Vebrugge and Garner (2009) illustrates how one can depart from this equilibrium. This model begins with a simple user cost known as the Verbrugge Variant:

$$u_t = P_t (i_t + \gamma_t - E\pi_t) \quad (3.5)$$

Where P_t represents house price, i_t represents interest expense, γ_t collects terms like depreciation, maintenance, insurance, property tax, other operating costs, and may or may not include a risk premium. $E\pi_t$ represents expected housing appreciation. All terms are in period t . From this simple framework, a more complex model is constructed by adding in the the preferential tax treatment homeowners receive:

$$u_t = P_t \left(i_t(1 - \tau_t^{Fed}) + \tau_t^{Prop}(1 - \tau_t^{Fed}) + \tilde{\gamma}_t - E\pi_t \right) \quad (3.6)$$

Where τ_t^{Fed} is the federal marginal income tax rate, τ_t^{Prop} is the property tax rate, and $\tilde{\gamma}_t$ is defined as before except now excluding property tax.

Importantly, this is specifically the user cost for owner-occupied homeowners, not landlords for whom taxes would differ significantly. The presence of differential tax treatments illustrates one reason why rents and user costs may diverge.

There are a variety of other theoretical reasons that might explain why user cost and rents would diverge. For example, landlords change rents infrequently (see Chapter 4).

Another explanation is the heterogeneous market for higher end homes. Diewert

and Nakamura argue that there are fewer homes for rent in the higher price range, and they are more likely to reflect house sitting arrangements including renter responsibilities. Moreover, given that housing is a matching market, and that both the supply and demand of higher end houses for rent are likely to be thin, it would be unsurprising if rents for similarly priced homes could vary greatly from market to market depending on the relative number of individuals looking to rent or be rented to.

The empirical literature on the equality of rents and user costs lends support to the inconclusive theory.³ Collectively, the literature shows that the inequality of rents and user costs holds at different levels of aggregation and is robust to a variety of specifications of theoretically justifiable user cost and rent measurements. Verbrugge (2008) produced a user cost with a closer relationship to rents, however this required the use of expected inflation as a proxy for appreciation, and this remains a theoretically suspect measure of user costs. In fact, as Gillingham argued, one could simply impose the condition that the proper appreciation or interest rate is the one which generates user cost equal to rents. However, given that this requires data on rents, it is more direct to simply use equivalent rents as the measure of housing services.

This research lead Verbrugge (2008) to conclude that rents were empirically not a good measure of user cost. In 2009 Diewert, Nakamura, and Nakamura declared, “The time has come, we feel, to accept the evidence of Verbrugge and others that user costs and rents do not reliably move together”. Importantly, this left price statisticians requiring a new justification for OER. As Verbrugge (2004) put it:

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 re-

³See Verbrugge (2004, 2008) and Verbrugge and Garner (2009), and Gyourko and Glaeser (2007).

jected. Statistical agencies must think about, and be able to justify, the continued use of rental equivalence on other grounds.

3.5 Does User Cost Justify Marginal or Average Rents?

The emergence of both theoretical and empirical evidence for the divergence has lead to a decline in the user cost justification, notably by Verbrugge and his co-authors at the BLS. Nevertheless, the user cost approach remains a useful theory by which OER is still sometimes judged. Deiwert, Nakamura, and Nakamura (2009), for instance, considered it to still be such an important underpinning that they argued when user cost and OER diverge, one should sometimes utilize user cost instead. Therefore, in attempting to determine whether marginal or average rents are appropriate for the CPI, it is useful to consider which is justified by user cost theory. To determine this, there are two important questions that must be answered: “Do current user costs equal marginal or average rents?” and “Are current user or average user costs the correct measure of housing services according to cost-of-living index theory?”.

The equality of rents and user costs comes from basic Jorgensian capital theory, and it says that rents and user costs of the same period should be equal. The universally reported conclusion in the literature is that $u_t = R_t$, with both user costs and rents sharing the same time subscript t . There is nothing in the theory which implies that user costs should equal an average of lagged rents. If anything, the evidence suggests that house prices respond more quickly than rents in response to changes in local housing amenities (Lang, 2012). This implies market rents are already a lagged measure of user costs, a problem that further lagging would exacerbate.

The literature produces some examples of agreement with the common sense claim that current user costs and average rents are not equal. For instance, Garner and Verbrugge (2009) argue:

...since market rents typically change on an annual basis, market rents inevitably reflect lagged conditions – and thus, perhaps, lagged rather than current user costs – potentially making it more difficult to discern a relationship between these variables.

However, given that they refer to rents changing on an annual basis, when they say “market rents”, they are referring to average rents, not marginal rents. This statement then endorses the view that average rents are not equal to current user costs. Verbrugge and Garner (2009) present more specific arguments in support of this claim:

A second issue related to the interest rate is that of marginal versus average user cost. A quarterly user cost measure will most naturally be a current user cost, i.e., it will incorporate the current period home price and the current period interest rate. However, rent indexes generally do not share this temporal feature. Instead, these indexes are averages constructed from a sample of all existing rent contracts, rather than from a sample of new contracts each period; thus, these indexes are implicitly temporally aggregated, being averages of contracts that were renewed this month, renewed last month, and so on.

Therefore, Garner and Verbrugge (2009) concur with the common sense conclusion from basic theory that if current user costs are desired, then the existing, average rents are not correct. Instead, the theory predicts that marginal rents are the appropriate proxy for current user costs.

The remaining question then is whether current user costs are what should be used. Fortunately, it is relatively straightforward application of cost-of-living index theory, which the BLS acknowledges as the measurement goal of the CPI, to see that current user costs are what the CPI should be measuring. Consider again the question that a cost-of-living index seeks to address:

What is the cost, at this month's market prices, of achieving the standard of living actually attained in the base period?

COLI theory is concerned with measuring changes in current costs, and since current user cost measures this month's cost of consuming housing services, they are more consistent with COLI goals of the CPI than some average user cost measures would be. Similarly, one can consult the basic formula for a COLI which compares utility as a function of prices in periods t and $t + 1$. Prices in periods prior to the base period play no role.

Therefore, given that the CPI attempts to measure current user costs, and that current user costs equal marginal rent, it implies that if one could construct a rental measure that was marginal rather than average rents, it would be preferable. As long as the user cost justification for rental equivalence is used, marginal and not average rents are justified.

3.6 Opportunity Cost

The equivalence between user cost of housing and market rents has long been used to justify the use of rental equivalence to measure inflation in OOH. Given this, the theoretical and empirical problems with the user cost theory discussed above could also be seen as an issue for rental equivalence. For this reason, economists have increasingly recognized the so-called "opportunity cost" justification for rental equivalence as being more theoretically important (Pool, Ptacek, Verbrugge, 2005; Diewert and Nakamura, 2009; Diewert, Nakamura, and Nakamura, 2009).

Poole, Ptacek, and Verbrugge frame the opportunity cost approach as asking:

How much richer would the homeowner be if he or she did not consume the housing services provided by a dwelling?

To answer this, it is argued, one should look at the rents that homeowners could have earned had they rented their homes instead of consuming the housing services.

Importantly, this is the same conclusion arrived at from user cost theory. However, in contrast to the user cost justification for rental equivalence, this justification requires no assumptions about the housing market being in equilibrium. For example, a transactions cost lead disconnect between the user cost and market rents for a given unit has no implications here. Regardless of the disequilibrium in the housing market, a homeowner still forgoes the opportunity cost of renting a unit by choosing to live there rather than rent it out.

This justification allows the cost of housing services for renters and homeowners to fit within a single framework. For renters, rental equivalence uses the actual rents, whereas for owners, it uses the implicit rents. However, the conceptual objective for the both is identical: “How much richer would the homeowner be if he or she did not consume the housing services provided by the dwelling?”.

One possible objection to this is that the opportunity cost of homeownership is not what a owner could have rented their home for because *they had to rent something*. This would imply that the opportunity cost is what they could have rented their house for less the cost of their next best rental option.

Consider, for example, a household that owns and lives in a house which they could rent out for R_1 , and if they did rent their house out, their next best option would be to rent a different housing unit for R_2 . One could argue, based on the fact that a household must consume housing, that the opportunity cost should be $(R_1 - R_2)$. However, to apply this concept consistently to the CPI would be to say that the opportunity cost of consuming a good is not the cost of that good, but the net cost relative to the next best option. However, as Poole, Ptacek, and Verbrugge (2005) correctly point out, all goods in the CPI have conceivable next best options that are ignored. The cost of a Red Delicious apple for the purposes of the CPI is not the price of that apple less the price of a Granny Smith, despite the fact that the consumer could have purchased the Granny Smith.

Cost-of-living theory clearly supports the notion that next best options should be ignored. Consider an individual consumes the same basket of goods in two periods, and none of the prices of those items change. A cost-of-living index in this example would clearly be unchanged regardless of what happens to the prices of the next best choices that the consumer did not consume. Even if Granny Smith apples are the consumer's next best choice, if an individual consumes none of them then their cost of living will be unaffected by changes in their price.

Another objection to the opportunity cost justification for rental equivalence is that the user cost can, under certain circumstances, be larger than the opportunity cost, and that the real cost should be the greater of the two (Diewert, Nakamura, and Nakamura, 2009; Deiwert, 2009; and Nakamura and Deiwert, 2009). It is argued that the opportunity cost that homeowners give up is not just the forgone rent from consuming their home, but the forgone returns from alternative financial investments. Deiwert (2009) frames the question like this:

Perhaps the correct opportunity cost of housing for an owner occupier is not his or her internal user cost but the maximum of the internal user cost, which is the financial opportunity cost of housing, and what the property could rent for on the rental market. After all, the concept of opportunity cost is supposed to represent the maximum sacrifice that one makes in order to consume or use some object.

Therefore it is argued that the true cost of housing services in a given period is whichever of these two opportunity costs is larger.

However, the financial investment aspect of housing represents housing as an asset, and is thus outside the scope of the CPI. Deiwert and Nakamura (2009) recognize this fact and argue the value of the flow of housing services is necessarily entangled with the asset nature of housing, because “there is no way of living in a home without investing in housing.” This certainly makes housing unique in the CPI, but it does not

explain why the investment component should be included if the two can be estimated separately.⁴ In addition, ignoring the financial opportunity cost is consistent with the BLS's treatment of other goods. Every good a consumer purchase has a financial opportunity cost since they could have chosen to invest the money rather than use it to purchase that good. As with other next-best purchase opportunity costs, financial opportunity costs should be ignored.

A third possible objection to the opportunity cost justification is related to constructing the proper counterfactual. One could object that in the counterfactual where all owner-occupied homeowners rented out their homes in a given period, the market for rental homes would be flooded with a large quantity of supply, and the price of rentals would go down.⁵ Therefore, the prices that would be observed in this counterfactual are quite different than current market rents. However, a counterfactual of all homeowners choosing to rent their homes out is irrelevant to the prices and choices any individual faces. An individual's opportunity cost only includes real and relevant choices, which are to either own or rent at current market prices.⁶

Despite the objections, opportunity cost remains the preferred justification of OER by the BLS and leading price statisticians. It has a clear advantage over the user cost

⁴There are goods in the CPI with asset values, including clothing and household durable goods. However, these goods typically do not include an investment component in that asset appreciation and is not part of the benefit of purchasing the goods.

⁵This assumes that those homeowners didn't then consume other housing, which would increase the demand at the same time.

⁶Another complication with the opportunity cost framework is the treatment of landlord costs. If a homeowner were actually to rent their home, then the benefit they would receive would be the rent less landlord costs. This would include things like vacancy costs, billing, the advertising costs of locating tenants, and other costs of business for landlords. In addition, the principal agent problem inherent in a landlord-tenant relationship suggests that maintenance costs will be higher for a rented unit (Henderson and Ionannides, 1983). These higher normal maintenance costs would have to be deducted from the rent as well. However, Deiwart (2003) has argued, in the context of user costs, that landlord specific costs should not be considered in-scope for CPI measurement and should be subtracted. Thus, perhaps the correct conceptual opportunity cost is one often embraced in OER literature: the opportunity cost to a homeowner of renting the unit to themselves. This avoids the principal agent problem and thus the additional ordinary maintenance costs. It also avoids many landlord specific costs of business since there is zero vacancy, tenants do not need to be billed, and advertising costs are unnecessary. However, since expected landlord costs are priced into market rents used to estimate OER, it is difficult to see how this theoretical approach justifies measurement goals that would be possible in practice.

justification in not requiring strong assumptions about housing market equilibrium, and given that user cost and rental equivalence in practice have not held up to empirical scrutiny. Ultimately, whatever measure of OER is used should be judged against this conceptual benchmark.

3.7 Does Opportunity Cost Justify Marginal or Average Rents?

As with the user cost, it must be determined whether the opportunity cost justification for owners' equivalent rent justifies marginal or average rents. Fortunately, the counterfactual choice that the opportunity cost considers makes it clear that what should be considered is market rents. If a homeowner were to rent her housing services the price she could get is the current market price for those services. The average price consists of current market prices from past periods which are irrelevant to someone who wishes to sell housing services in the current market. The average price is a choice which is unavailable in the current period.

The focus on current prices is also clear from COLI theory. For example the BLS Handbook states:

As it pertains to the CPI, the COLI for the current month is based on the answer to the following question: What is the cost, at this month's market prices, of achieving the standard of living actually attained in the base period?

Cost of living index measures the change in welfare from one period to the next, and thus focuses on the standard of living attainable in each period. Therefore, what is relevant for the opportunity cost of housing is the change in price of a choice not taken but necessarily available. The available prices in periods t and $t + 1$ are the prices a homeowner could have rented their home for in period t , and the price they could have rented it for in period $t + 1$. In period t , the average rental price is not available for a homeowner to choose in this period, nor is the average price in $t + 1$

available then. In each period, only the current market prices, and thus marginal prices, are available and therefore relevant.

There is some empirical evidence that homeowners also consider market rents and not average rents to be the value of their opportunity costs. Hoffman and Kurz (2004) look at estimates of implicit rent made by homeowners in a national survey from Germany. This survey asks homeowners:

And if you lived in this flat or house as a tenant: what do you estimate would be the monthly rent without heating costs?

Note that the framing of the question biases the results in favor of an average rent. In contrast, a framing that would more accurately reflect the BLS's stated goal of the CPI would explicitly mention current prices. Despite this, the measure of implicit rent estimated by homeowners tracks marginal rents in the German CPI more closely than average rents. Figure 3.1 below, from Hoffman and Kurz (2004), shows that average rents diverge from rental equivalents reported by homeowners, whereas rents in new contracts, which reflect market rents, track them more closely.

Figure 3.1: Rent Measure Comparison
Source: Hoffman and Kurz (2004)

While the average rents and implicit rents are both smoother, the average rents have large deviations from new contract rents. From 1985 to 1992, rental equivalents

rose 54%, and new contract rents rose 42%. In contrast, rents on average rose a significantly smaller 28%.⁷

In addition, in one of the few papers to explicitly consider the issue, Shimizu, Deiwert, Nishimura, and Watanabe (2012) conclude that:

Conceptually, the imputed rent is a rent level that a house owner can receive when leasing the house in the rental house market today.

Therefore, the imputed rent always matches the market price.

Overall, the underlying theory, empirical evidence, and the general conclusions of research in this literature clearly suggest that the opportunity cost justification for OER implies the CPI should measure the changes in marginal and not average prices.

⁷Data for these calculations provided by Claudia Kurz-Kim.

CHAPTER 4

STICKY RENTS AND TENURE DISCOUNTS

As discussed in the beginning of Chapter 3, the CPI should be measuring marginal rather than average rents. For most goods, this measurement difference would not be very consequential. For non-durable goods, the average price, as measured by a survey of household spending, should be close to marginal prices of the actual goods and services consumed in that period, as measured by changes in average market prices.¹ In contrast, rents are sticky, meaning they have nominal price rigidity. In large part due to long-term leases, rents that households pay every month change very slowly, and respond slowly to changes in market prices as contracts expire and are renegotiated. In contrast to other goods where marginal and average prices should be close, this by itself can create a large divergence for rents. The following chapter will provide an overview of sticky rents, and the related phenomenon of tenure discounts.

4.1 Sticky Rents

The primary cause of sticky rents is that landlords and tenants typically agree to leases that are a year or more in length. The 1995 Property Owners and Managers Survey (POMS) supplement to the AHS reported that 85% of tenants have a lease, 43% of which are for a year, and 2.3% are for more than a year.

However, even after accounting for the long nature of contract length, rents often still do not change. Evidence on nominal rigidity in U.S. rents can be found in Genesove (2003), who documents rental rigidity in the U.S. from 1974-1981, and Verbrugge and Gallin (2012), who characterize rental rigidity using BLS micro data from 1998-2011.

¹Differences can arise if, for example, households disproportionately purchase items that are on sale, but this is the extent of the wedge that can be driven between market and average non-durable good prices. There is evidence of sticky prices for non-durables, but they are usually found in small value goods, whereas housing constitutes 20% to 30% of a households budget (Genesove, 2003).

Table 4.1: Average Lease Length

Lease Length	Percent
No Lease	15.5%
Less than one year	36.1%
Annual	44.4%
Greater than one year	4.0%
<i>Source: POMS; Crone, Nakamura, Voith (2011)</i>	

Genesove creates an annual panel of AHS data from 1974 to 1981 which is unique because in all other periods the national AHS survey occurs only every other year. Annual change in rent is measured by looking at percent change in rent for each unit from one year of AHS data to the next. In all years there is strong evidence of nominal rigidity, with the percent of units with zero change in rent in a given year ranging from 23% to 34%, and averaging 29% across all seven years. He also reports on nominal rigidity in BLS micro data from 1988 through 1992, which shows that 37% of units surveyed had zero change in rent over 18 month periods. Verbrugge and Gallin find a larger amount of rigidity for 1998-2011, with 54% of units experiencing no change after 12 months. At 18 months they find 41% experience no change, which is higher than but close to what Genesove found for 1988 through 1992 with the same BLS micro-data. Collectively this evidence indicates that nominal rigidity has increased over the previous 37 years.

The rise in nominal rigidity is likely related to Genesove's finding that the faster the underlying rate of rental inflation, the less likely that nominal rigidity would be observed. He found that a 1% increase in median rent growth decreases the incidence of nominal rigidity by 2.2%. This result is what would be expected, as faster growth in market rates makes nominal rigidity more expensive for landlords. From 1974-1981, the average 12-month change in the CPI for rents was 6.7%, from 1988-1992 it fell to 3.6%, and from 1998-2011 it fell further to 3.1%. Given this decline in the underlying inflation rate, it is unsurprising that nominal rigidity has increased.

In addition to nominal rigidity at 12-month changes, Verbrugge and Gallin found that for a significant portion of units there was no change in rent at longer intervals as well. Table 4.2 summarizes. This shows that nominal rigidity remains even after two

Table 4.2: Probability of Rent Change

Duration	Percent Decrease	No Change	Percent Increase
6-month	9%	69%	22%
12-month	11%	54%	36%
18-month	14%	41%	45%
24-month	15%	34%	51%

Source: Verbrugge and Gallin, (2012)

years for 34% of units. While there are some decreases, increases are around three times more common. Using the Kaplan-Meier product-limit estimator, Vebrugge and Gallin estimate hazard functions for rent survival time, with survival defined as not undergoing a change in rent. They find that the odds of a rent change in given period, conditional on no change occurring up that period, generally fall over time. The probability of a rent change is generally halved around the 3rd or 4th year of tenure.

There are a variety of factors that are related to the odds of nominal rigidity. One important factor appears to be structure type. Genesove found that single family units experience nominal rigidity after twelve months half of the time, compared to 13% for apartment buildings with 50 or more units. Verbrugge and Gallin report single family detached units have 68.2% rigidity, compared to 47.1% for multifamily units without an elevator and 33.7% for multifamily units with an elevator. Similarly, he finds that units with continuing tenants experience systematically lower rent growth even when the growth is greater than zero.

In the international context, Hoffman and Kurz-Kim (2006) provide evidence of nominal rent rigidity for Germany. They found that around 2.2% of rents experience a price change each month, corresponding to 1.5% of rents increasing, and 0.7%

decreasing. They found that 30% of (quality adjusted) rents did not change at all in the six years of data they looked at. The average size of price changes was 9.8% for price increases, and -9.2% for price decreases. For Japan, Shimizu, Nishimura, and Watanabe (2008) provide evidence of nominal rigidity. The probability that any unit experiences no change in price in a given year is 89.3%.

The evidence presented thus far demonstrates that a large share of rent changes are zero percent. However, Hoffman and Kurz-Kim (2006) argue that nominal rigidity cannot be demonstrated simply by observing the average frequency and sizes of price changes. Few price changes may be observed because nominal prices are rigid, or because prices have little reason to change. Low volatility in input costs, or demand, or other structural factors could lead to very little price change even with perfectly flexible prices. However, there is no obvious reason why landlords would change rents at exactly zero rather than just a very small price change. The volatility of house prices does not suggest stable demand for housing or opportunity costs given that landlords have the option of selling their buildings at volatile market prices. Without a plausible explanation for why prices have little reason to change, and given the underlying volatility of housing, the high frequency of zero price changes should be interpreted as demonstrating nominal rigidity.

Given the compelling evidence of nominal rigidity in rents, the natural question is, “Why does it occur?”. In the literature on nominal rigidity in other goods, menu costs are a common explanation for sticky prices. However, landlords have no literal menus to change. Furthermore, while menu costs typically apply to a single seller of a homogeneous good that sets one price for many customers, housing is a heterogeneous good where prices are usually negotiated by bargaining between the two parties (Genesove, 2003). Therefore, the typical explanation for rent stickiness does not apply, and a different explanation must be found.

One important explanation is grid pricing, when rents tend to change in intervals

of \$50, \$100, or other discrete intervals greater than \$1. Verbrugge and Gallin (2012) document extensive grid pricing in U.S. rents, with 25% of rents being multiples of \$100, and 92% of rents are multiples of \$5. Using probit regression, Genesove (2003) estimates that half of the nominal rigidity can be attributed to grid pricing. However, this still leaves a substantial amount of rigidity to be explained. In addition, grid pricing still represents nominal rigidity given that optimal prices at any given time are unlikely to be exactly round numbers (Verbrugge and Gallin, 2012).

4.2 Tenure Discounts

After grid pricing and the long-term nature of leases, the remaining explanations for nominal rigidity can generally be regarded as falling under the category of tenure discounts. There are a variety of potential causes for these discounts, but the shared feature is that tenants receive rents below market value as their tenure length increases. Gausch and Marshall (1983) distinguish between a sit discount, which is a discount received by tenants only at a first contract renewal, and a length-of-stay discount, which increases with tenure length.

There are several studies providing empirical evidence on the extent of tenure discounts. Goodman and Kawai (1985) use hedonic regression on 1977 AHS data rents to estimate tenure discounts for 19 cities. They find consistent evidence of discounts in all cities. The percentage discount for each year of occupancy ranges from -0.52% in Salt Lake City to -2.46% in Spokane. The average discount is -1.29% per year. Looking at the average discounts within cities, which illustrate you how much rents would increase if current rents had no tenure discount, they find a range from 1.1% in Salt Lake City to 8.2% in Boston. Malpezzi and Follain use 1974-1976 AHS data for 39 cities and find an average discount of 1% for continuing tenants. This is statistically significant and negative in all 39 cities.

In addition to hedonic regression approaches, one can compare nominal rigidity

in units that receive a new tenant to units that have a continuing tenants. Genesove (2003) shows that 36% of units with continuous tenants experience nominal rigidity, compared to 14% for units with new tenants. In addition, he found that the median growth rates were higher for units with new tenants than for units with continuing tenants. In earlier work using the same data, he showed that third year tenants had rents that were an additional 5% lower (Genesove, 1999). Therefore, by the end of the third year, tenants had rents that were 9% below market rate. As shown in table 4.3, Verbrugge and Gallin (2012) found that 12-month rent changes, both positive and negative, are more common when a new tenant moves in. Some continuing tenants appear to receive discounts via lower odds of an increasing rent, and some new tenants appear to receive discounts via higher odds of a rent decrease. However, of the 24.9% fewer units with nominal rigidity among new tenants, 15.4% have price increases compared to 9.6% having price decreases. This may reflect a greater prevalence of discounts for continuing tenants, or it may reflect higher nominal rigidity for continuing tenants and rents simply being more likely to increase over time than decrease.

Table 4.3: Probability of Twelve-Month Rent Change by Tenure

	Percent Decrease	No Change	Percent Increase
Continuing Tenant	10.1%	52.5%	37.4%
New Tenant	19.7%	27.6%	52.8%

Source: Verbrugge and Gallin (2012)

Alternatively, the higher prevalence of rent decreases among new tenants may reflect a discount for new tenants. Barker (2003) presents a model where long-term tenants should be charged more than short-term tenants because their price elasticity is lower. Depending on turnover costs, this might mean overall new tenants should be charged less. His survey of 100 apartment managers found that first month discounts are more prevalent than renewal discounts, suggesting there is some truth to this.

While this may be true for a subset of units, and may partly explain the higher prevalence of decreases among new units found in Verbrugge and Gallin, the evidence suggests that on average, tenure discounts exist and are more common than new tenant discounts.

In the international context, Shimizu, Nishimura, and Watanabe (2008) found nominal rigidity is higher for units with continuing tenants renewing their contracts, who have a 97.0% chance of no price change, than for units that get new tenants, who have a 75.5% probability. As with the U.S. case, the price nominal rigidity is asymmetrically downward. For units with existing tenants who are renewing contracts, there were no price increases, only decreases. In contrast, for units with new tenants, price increases made up around a quarter of the price changes.

Hoffman and Kurz-Kim (2002) find strong evidence of tenure discounts from hedonic regressions on rents from a survey of West German households. In the regressions, they control for tenure discounts and estimate an average tenure discount effect, shown in Figure 4.1 below.

Figure 4.1: Level of Rents by Tenancy Duration
Source: Hoffman and Kurz (2002)

While tenure discounts may explain why nominal rigidity occurs, an explanation for why tenure discounts occur is needed. Explanations given in the literature include transaction costs of renegotiations, transaction costs due to outside options, and random error plus censoring.

Genesove (2003) hypothesizes that the convention of using previous period's rent saves the cost of information acquisition, bargaining time, and emotional stress of renegotiation for landlord and tenant. Each of these could be included under the sub-category of "transaction costs of renegotiation."

A similar transaction cost based explanation focuses on the transaction cost of outside options. This occurs if, as tenure length increases, the cost of choosing an outside option increases for the tenant or the landlord. An example would be if the tenant-landlord relationship generates a surplus over time, perhaps due to moving costs on the part of the tenant, or new tenant costs on the part of the landlord (Genesove, 2003). There is evidence that new tenant costs on the part of the landlord are substantial with one estimate of turnover costs related to painting, decorating, and lost rents of around \$1,174 (Barker, 2003). Alternatively, a surplus may arise from landlord and tenant heterogeneity that, due to asymmetrical information, is unobservable ex ante but revealed ex post. This could allow matches between "good" tenants and "good" landlords to result in a surplus that can be split (Hubert, 1995).

Regardless of the causes, nominal rigidity of rents and tenure discounts are a widespread phenomenon. Importantly, nominal rigidity appears especially likely to affect the subsample of rents that receive the most weight in the CPI housing sample in the OER estimation: single family homes. Almost half of detached homes in Genesove's sample report nominal rigidity compared to 13% in apartment buildings with more than 50 units.

CHAPTER 5

IMPLICATIONS OF CURRENT OER MEASUREMENT

In Chapter 3, it was argued that the CPI should be measuring market rather than average rents. Chapter 4 presented evidence that rents exhibit nominal rigidity which will cause a divergence between a given unit's marginal and average rent. These results raise questions about using average rents instead of marginal rents and how it affects aggregate measures of rental inflation. This chapter will present the following conclusions using a simple model of rent:

- Average rents are biased downward in levels relative to market rents.
- A shock affects average rents with a lag compared to market rents.
- Over time, the effect of the shock on average rents will not converge to its effect on marginal rents unless long run price growth is zero.
- Price relatives are not biased in the long-run.
- The effect of a shock on price relatives using average rent will be biased towards zero in the short-run.

5.1 Effect of Nominal Rigidity on Rent Levels

The question of average versus marginal rents is, to some extent, an issue of whether and to what extent micro-rigidity translates to macro-rigidity. Micro-rigidity means that at the individual unit level prices are slow to adjust to changes in market prices, whereas macro-rigidity refers to the aggregate measure of prices adjusting slowly to changes in underlying market prices. A simple model illustrates how this macro-rigidity can occur.

Market rents R_t^* are assumed to follow the process:

$$R_t^* = R_{t-1}^* \delta_t \quad (5.1)$$

Where δ_t is the period t market rent growth rate. This implies that market rents in a given period can be written purely as a function of t past rents and beginning period rent R_0 .

$$R_t^* = \prod_{i=0}^t \delta_{t-i} R_0^* \quad (5.2)$$

A basic Calvo model of price determination will be used for individual rents, where in this case long-term leases are the only factor that generates nominal rigidity. In each period, a unit either receives a new lease, or the old lease remains in force. If the old lease remains in force, then there is no change in price, and the rents are equal to the previous period's rent. If a unit receives a new lease, the rent is equal to the market rate for that unit. The random variable θ is the proportion of units with the old lease in effect, and therefore $(1 - \theta)$ is the proportion of units with a new lease. As before, let R_t^* be the actual market rent in period t , the average rents are defined as:

$$R_t = (1 - \theta) R_t^* + \theta R_{t-1} \quad (5.3)$$

Using this model, it can be shown that average rents are less than market rents if long-run price growth is positive, but greater than market rents if price growth is negative. With the simplifying assumption that the growth in every period is equal to a constant δ , average rents can be solved iteratively so that R_t can be written as

a function of the growth rate δ , an initial market rent, and θ :

$$\begin{aligned} R_t &= \sum_{i=0}^t (1 - \theta) \theta^i \delta^{t-i} R_0^* \\ &= (1 - \theta) \delta^t R_0^* \sum_{i=0}^t \left(\frac{\theta}{\delta} \right)^i \end{aligned} \tag{5.4}$$

So long as $\delta > \theta$, the long-run rent can be derived by taking the limit of R_t as $t \rightarrow \infty$:

$$\begin{aligned} \lim_{t \rightarrow \infty} R_t &= \lim_{t \rightarrow \infty} (1 - \theta) \delta^t R_0^* \sum_{i=0}^t \left(\frac{\theta}{\delta} \right)^i \\ &= \frac{(1 - \theta)}{(1 - \frac{\theta}{\delta})} \delta^t R_0^* \end{aligned} \tag{5.5}$$

This is less than market rent $\delta^t R_0^*$ when the growth rate δ is greater than one, greater than market rent when δ is less than one, and equal to the market rent when δ is equal to one. Thus if rents are increasing over time, as is generally expected, average rents will be biased downward from market rents. Consider the illustrative values of 91.7% for θ , and 1.0025 for δ , which is a monthly growth rate that corresponds to an annualized growth rate of 3%.¹ Under these assumptions, equation 5.5 implies average rent will be equal to 97.3% of market rents. Under a faster annualized growth rate of 10%, average rents would be 88.7% of market rents.

5.2 Effect of Nominal Rigidity on Rent Changes

While average rents are biased in levels, the effect of shocks to market rents on average rent can be shown by looking at changes in the period t market growth rate δ_t . Again, let growth rates be equal to constant δ in all periods but t . Then all market

¹A θ of 91.7% means a unit has an 8.3% chance of receiving a lease in a given period, which translates to an average lease length of $1/8.3\% = 12$ months.

rents after period t can be written as:

$$\begin{aligned} R_{t+s}^* &= \delta^s R_t^* \\ &= \delta^s \delta_t R_{t-1}^* \quad \forall s \geq 0 \end{aligned} \tag{5.6}$$

Therefore, in market rents, changes in the growth rate have a permanent impact on future rents which are fully reflected in period t :

$$\frac{\partial R_{t+s}^*}{\partial \delta_t} = \delta^s R_{t-1}^* \quad \forall s \geq 0 \tag{5.7}$$

Next, the effect of shocks on market rents can be compared to the effect on average rents to determine if the impact is equal. Solving equation 5.6 iteratively shows average rents in period $t + s$ to be a function of past growth rates and θ .

$$\begin{aligned} R_{t+s} &= (1 - \theta) R_{t+s}^* + \theta R_{t+s-1} \\ &= (1 - \theta) \delta^s \delta_t R_{t-1}^* + \theta R_{t+s-1} \\ &= (1 - \theta) \delta^s \delta_t R_{t-1}^* + \theta \left((1 - \theta) R_{t+s-1}^* + \theta R_{t+s-2} \right) \\ &= (1 - \theta) \delta^s \delta_t R_{t-1}^* \sum_{i=0}^s \left(\frac{\theta}{\delta} \right)^i + (1 - \theta) \theta^{s+1} R_{t-1} \end{aligned} \tag{5.8}$$

This formulation conveniently divides the rent in period $t + s$ into two parts: the portion that is a function of δ_t on the left, and the remaining part which is not a function of δ_t contained on the right in $(1 - \theta) \theta^{s+1} R_{t-1}$. This allows the change in average rents in response to underlying inflation to be derived as:

$$\frac{\partial R_{t+s}}{\partial \delta_t} = (1 - \theta) \delta^s R_{t-1}^* \sum_{i=0}^s \left(\frac{\theta}{\delta} \right)^i \tag{5.9}$$

It is clear that in the first period, where $s = 0$, the impact on average rents of a shock to growth is less than the impact on marginal rents. Marginal rents change by δR_{t-1}^* ,

and average rents change by $(1 - \theta)\delta R_{t-1}^*$.

The impact in period $t + s$ of a onetime change in growth in period t converges over time, but not fully. So long as $\delta > \theta$, it can be shown that:

$$\begin{aligned}\lim_{s \rightarrow \infty} \frac{\partial R_{t+s}}{\partial \delta_t} &= \lim_{s \rightarrow \infty} (1 - \theta)\delta^s R_{t-1}^* \sum_{i=0}^s \left(\frac{\theta}{\delta}\right)^i \\ &= \frac{(1 - \theta)}{\left(1 - \frac{\theta}{\delta}\right)} \delta^s R_{t-1}^*\end{aligned}\tag{5.10}$$

From this, we can see that whether or not the long-run impact on average rents is less than, greater than, or equal to the impact on marginal rents depends on the constant growth rate δ . As long as the long-run growth rate is positive, e.g. $\delta > 1$, the impact of a one period shock on average rents is less than the impact on market rents. The closer the long-run growth rate is to 1, the more likely that the impact in a given period converges over time. In addition, given that for $\delta > 1$ equation 5.10 is strictly increasing in s and that the impact on average rents is less than on marginal rents in periods $s = 0$ and in the limit, it can be seen that the impact on average rents is less than the impact on marginal rents in periods $s = 1, 2, \dots$

5.3 Effect of Nominal Rigidity on the CPI Measure of Inflation

This result implies that, unlike marginal rents, average rents reflect changes in underlying market rates with a lag. However, as discussed in Chapter 2, the CPI measure of inflation is based on price relatives. The following equation shows a simplified² model of the price relatives that underly owners' equivalent rent:

$$PR_{t,t-1} = \frac{\sum R_{it}}{\sum R_{it-1}}\tag{5.11}$$

By dropping the i subscript and focusing on rent aggregates, the long-run price relative can be derived. Using equation 5.5 for $\lim_{t \rightarrow \infty} R_t$ and R_{t-1} , and again assuming

²This model ignores the lagging induced by sampling individual units every six months.

a constant growth rate δ , the price relative can be shown to be equal to the market growth rate:

$$\begin{aligned}\lim_{t \rightarrow \infty} \frac{R_{it}}{R_{it-1}} &= \frac{\frac{(1-\theta)}{(1-\frac{\theta}{\delta})} \delta^t R_0^*}{\frac{(1-\theta)}{(1-\frac{\theta}{\delta})} \delta^{t-1} R_0^*} \\ &= \delta\end{aligned}\tag{5.12}$$

Importantly, this shows that the bias in the levels does not imply that the price relatives are biased. Equation 5.9 shows the impact of a shock in period t to be a downward bias in period t average rents relative to market rents. This means the period t price relative will have unbiased $t-1$ rents in the denominator and downward biased period t rents in the numerator, and therefore will be downward biased overall. However, as equation 5.9 shows, the shock in period t also downward biases rents in $t+1$. This implies that period $t+1$'s price relative will have downward biased price levels in both numerator and denominator and therefore will not necessarily be downward biased overall. In fact, given the convergence in Figure 5.1 below, price relatives after the first period are likely to be upward biased, reflecting catch-up inflation as the shock moves through the system.

While deriving theoretical values for this bias is complex, an illustrative example can be provided. Consider the same assumptions above and an initial market rent of \$100, and the initial implied average rent of $97.3\% \times \$100$. In addition, in period $t=0$ let the growth rates experience a one-time shock, moving from an annualized growth rate of 3% to an annualized growth rate of 10%, which translates to an increase of 7% at an annualized rate or a 0.0055 change in the monthly growth rate. Several of the results found above are confirmed in Figure 5.1, which shows the ratio of the marginal and average rents. This illustrates that in the long-run the shock does not permanently affect the relative levels of average and marginal rents, due to catch-up inflation in the average rents.

Figure 5.1: Ratio of Change in Rent Measures
From 7% Change in Period 0 Annualized Monthly Growth

The next figure shows the same scenario except using the CPI measure of inflation, e.g. price relatives.³ Marginal rent inflation increases sharply in the affected period, then returns to the equilibrium annualized monthly growth rate of 3%. In contrast, average rent inflation increases slowly and then slowly converges towards the equilibrium 3% growth rate. This suggests that price relatives based on average rent reflect market changes slowly, and this biases CPI towards one during the shock and then away from one after the shock. Marginal rents therefore are likely to be more volatile than average rents. Divergence from marginal rents by average rents is in contradiction to the goal of the CPI of measuring current costs. Given the large weight of OER in the CPI, this divergence is likely to be significant. In addition, the slow convergence towards marginal rent also verifies the result that average rent price relatives are not biased in the long-run, unlike average rent levels.

³Again, this ignores the issue of six month lags built into the CPI measure and uses one month price relatives rather than the sixth root of six month price relatives.

Figure 5.2: Annualized Price Relatives With 7% Change in Period 0 Annualized Monthly Growth

5.4 Changes in Turnover

The previous sections made the simplifying assumption that average turnover remained constant over time, and so θ was constant. Turnover does not affect marginal rent, therefore it should not affect inflation if it is properly measured. However, changes over time in turnover, or equivalently changes in average tenure length, are another reason why average and marginal rents may diverge. This is an especially salient concern given the current housing market trend of single-family housing units being converted to rental units, and the increasing proportion of households renting. In addition, it seems plausible that turnover may be subject to seasonal variation, which means turnover led variation may increase the seasonality of the CPI. Consider the following extension of the simplified model of average rents which allows θ to vary over time:

$$R_t = (1 - \theta_t) R_t^* + \theta_t R_{t-1} \quad (5.13)$$

Taking the derivative of R_t with respect to θ_t , it can be shown that:

$$\frac{\partial R_t}{\partial \theta_t} = R_{t-1} - R_t^* \quad (5.14)$$

This means that if market rent is above average rent, as the previous section shows is implied under the assumptions of constant turnover and positive growth rate δ , then decreases in θ_t , meaning faster turnover, increase average rent in period t .

In equilibrium, and given a constant θ up until point t , the R_{t-1} in equation 5.13 can be substituted out for a function of market rents R_{t-1}^* :

$$\begin{aligned} R_{t-1} - R_t^* &= \frac{(1 - \theta)}{(1 - \frac{\theta}{\delta})} R_{t-1}^* - R_t^* \\ &= \frac{(1 - \theta)}{(1 - \frac{\theta}{\delta})} \delta^{-1} R_t^* - R_t^* \\ &= R_t^* \left(\frac{(1 - \theta)}{(1 - \frac{\theta}{\delta})} \delta^{-1} - 1 \right) \end{aligned} \quad (5.15)$$

This implies that a decline in θ_t (meaning an decrease in average tenure length and an increase in turnover) will increase average rents so long as $\delta > 1$, and vice versa. Because market rents are not impacted by θ , this will lead to a divergence in changes of market rents and average rents, as well as a short-run bias in price relatives. As an illustrative example, let annualized rent growth equal 3%, $\theta = 91.7\%$, and period t turnover increase by one percentage point so that $\theta_t = 90.7\%$. This implies a positive shock to average rents equal to .03% of the period t market rent, despite no change in market rents.

By solving R_{t+s} iteratively to be a function of θ_t , the impact of a period t shock to turnover in later periods can similarly be shown to be:

$$\frac{\partial R_{t+s}}{\partial \theta_t} = \theta^s (R_{t-1} - R_t^*) \quad (5.16)$$

The impact of a change in turnover on the CPI measure of inflation can be found by taking the derivative of the period t price relative with respect to θ_t :

$$\begin{aligned}
\frac{\partial}{\partial \theta_t} \frac{R_t}{R_{t-1}} &= \frac{\partial}{\partial \theta_t} \frac{(1 - \theta_t) R_t^* + \theta_t R_{t-1}}{R_{t-1}} \\
&= 1 - \frac{R_t^*}{R_{t-1}} \\
&= 1 - \frac{(1 - \frac{\theta}{\delta})}{(1 - \theta)} \delta
\end{aligned} \tag{5.17}$$

A decrease in θ_t increases the measured inflation rate so long as δ is greater than one. Using the same illustrative values as above, the period t inflation rate increases by .03%, or 0.36% at an annualized rate, in response to the 1 percentage point decrease in θ_t . Higher turnover leads to higher measured inflation if the growth rate is positive and lower measured inflation if the growth rate is negative. To take a simple case where this is clearly true, consider if monthly turnover increased from 10% to 100% (meaning θ goes from 90% to 0%). Assuming δ is positive, which means rents are increasing overall, average rent would increase from below marginal rent to equal to marginal rent.

While single-period shocks are important to examine, longer term gradual shifts in turnover are another possible concern. This can be simulated using the calvo model of rents and some illustrative assumptions as before. A long-run annualized growth rate of 3% is again assumed and so is a starting monthly turnover rate of 8.3%, so that $\theta = 91.7\%$. The turnover rate increases by 0.01 percentage point per month, so that after 10 months, turnover has permanently increased by 1 percentage point to 9.3%, and therefore $\theta = 90.7\%$. This would represent a decrease in average lease length from 12 months to 10.8 months. Figure 5.3 below shows that the price relative based on average rents increases gradually, whereas as the price relative based on marginal rents is unaffected.

Figure 5.3: Annualized Change in Price Relatives
From Gradual Increase in Turnover

As tenancy gradually increases, so does the inflation rate, rising from .25% to .27%, which corresponds to a change in annualized inflation from 3% to just under 3.3%. The average inflation rate converges back towards the marginal inflation rate, but this process is still occurring more than two years after the turnover rate stopped changing. This implies that, unlike changes in growth rates modeled above, changes in turnover have a permanent impact on relative inflation levels of average and marginal rents. Figure 5.4 below confirms that the ratio of marginal to average rent levels is permanently reduced.

Overall, this shows that changes in the turnover rate can have important short and long-term impacts on inflation measured using average rent.

5.5 Endogenous Price Changes

The previous section showed that exogenous changes in turnover could lead to changes in the inflation rate. A further extension of the model shows that endogenous turnover can impact inflation as well and may mitigate macro-rigidity in the presence of micro-rigidity.

Figure 5.4: Ratio of Marginal to Average Rent
From Gradual Increase in Turnover

Caballero and Engel (2007) showed that aggregate price levels may be more flexible than is implied by nominal rigidity at the micro level. Shimizu, Nishimura, and Watanabe (2008) (hereafter “SNW”) and Verbrugge and Gallin (2012) both apply the model of Caballero and Engel to housing rents and investigate the extent to which micro-rigidity translates into macro-rigidity. SNW utilize a panel dataset of rents from Japan, and Verbrugge and Gallin use BLS housing micro-data from 1998 to 2011. Importantly, while the previous sections defined changes in rent as changes in leases, this model emphasizes the fact that not all lease changes are rent changes, and not all rent changes are lease changes. Rather than being a function of lease lengths that are determined in the past, this model considers rent changes to be a function of current and past rental inflation rates.

SNW define R_{it}^* as the rents for unit i in period t absent any micro-rigidity, but it could also be thought of as market rents. Changes in these rents are modeled as a function of aggregate shocks ξ_t and idiosyncratic zero-mean unit level shocks v_{it} :

$$\Delta \ln R_{it}^* = \Delta \xi_t + v_{it} \quad (5.18)$$

For each unit, the gap between the actual rent R_{it} and the market rent R_{it}^* is defined as:

$$X_{it} = \ln R_{it} - \ln R_{it}^* \quad (5.19)$$

The probability of a rent change is modeled as a function of X_{it} .

$$\Lambda(x) \equiv P(\Delta \ln R_{it} \neq 0 | X_{it} = x) \quad (5.20)$$

This model of micro-rigidity can then be used to illustrate how the measured average rent changes in response to aggregate shocks. Define the change in rent in period t for unit i given an aggregate shock ξ_t to be $\Delta \ln R_{it}(\Delta \xi_t, X_{it})$, and the cross section distribution of the gap between actual and market rents to be $h(x)$. Then the average change in rents can be represented as:

$$\ln R_{it}(\Delta \xi_t, X_{it}) \equiv \int \ln R_{it}(\Delta \xi_t, x) h(x) dx = - \int (x - \Delta \xi_t) \Lambda(x - \Delta \xi_t) h(x) dx \quad (5.21)$$

This equation shows that the change in average rents is a function of the size of the gap between actual and market rents plus the value of the shock, the probability that actual rents change, and the distribution of the gap between actual and market rents.

The response of average rents to a shock can be seen by deriving an impulse response function, which is Caballero and Engel's measure of price flexibility:

$$\lim_{\Delta \xi_t \rightarrow 0} \frac{\Delta \ln R_t}{\Delta \xi_t} = \int \Delta(x) h(x) dx + \int x \Delta'(x) h(x) dx \quad (5.22)$$

If there were no macro-rigidity, then this would be equal to one, as all shocks would be fully reflected in average rent changes. The first term on the right hand side reflects a result from the previous section: the more frequent rent adjustments are, the more responsive average rent will be to shocks, and thus the lower macro-rigidity will be. This intensive margin is the more obvious and direct way in which micro-

rigidity leads to macro-rigidity. The second term, the extensive margin, shows that macro-rigidity also depends on the change in the frequency of rent adjustments that a shock causes. If this number is positive, then the impulse response could be equal to one, reflecting zero macro-rigidity even if relatively few rents change, e.g. even if $\Delta(x)$ is low on average.

To illustrate how this can occur, consider the example of when market rents are above average rents, so that x is negative. In addition, assume that the more negative this gap is, the more likely landlords are to change rents, meaning that $\Delta'(x)$ is negative (because the more negative x , the higher $\Delta(x)$). In this example, it is clear that the impulse response function will be higher than would otherwise be the case. Importantly, this could be a negative or positive number, so that the extensive margin effect may increase or decrease macro-rigidity.

Verbrugge and Gallin (2012) decompose rent changes into the extensive and intensive margins and find that extensive margins remain relatively stable, whereas intensive margins move with the overall growth rate of rents. Furthermore, they find that even after controlling for spell length, units with a larger price gap are more likely to experience a rent change. Therefore, macro-rigidity is likely to be less than is implied by micro-rigidity.

SNW provide empirical evidence that the effect of extensive margins in Japan is minor. They estimate the Caballero-Engel measure of price flexibility to be 0.0097 indicating a high degree of macro rigidity. To illustrate the impact of micro-rigidity on inflation, they simulate how overall CPI for Tokyo, Japan would have looked under alternative levels of nominal rigidity. From their micro data, they estimate the actual quarterly frequency rent changes for Japan in the period under consideration was 2.5%. They then use rent change frequencies that correspond to 1) the level found in Germany from Hoffman and Kurz-Kim (2006), 2) the level found in the U.S. in the 1970s from Genesove (2003), and 3) a level in between these two. A graph of their

results is shown in Figure 5.5.

Figure 5.5: Simulations of Tokyo CPI Under Alternative Assumptions
Source: Shimizu, Nishimura, Watanabe (2008)

Overall, they found that using German level rigidity did not change the CPI significantly, but the intermediate rigidity between the U.S. and German did. Furthermore, using U.S. level of rigidity, they found that inflation increased by as much as 1% faster during the house price bubble period in Japan, decreased by as much as 2% faster during the house price bust, and showed deflation starting a year earlier. In contrast, during the relatively stable mild deflationary period at the end, there was far less divergence between the two measures. These are significant changes in the inflation rates. In addition, they do not investigate what the divergence would be if fully flexible rents were used, and instead looked at rent flexibility corresponding to U.S. and German levels. If the actual market rates were used, as advocated in this paper, instead of a weighting of market and previous period rents, then the differences would have been even greater.

5.6 Additional Lags in Current CPI Measurement

The above sections suggest that the current CPI measurement will reflect changes in underlying market rents only with a lag due to the use of average rents. In addition, the current BLS measurement procedure introduces an additional measurement lag by using the sixth root of 6 month changes in rent to estimate monthly inflation. This induces a three month lag, which makes the CPI less timely than if monthly changes in rents were used. The BLS has traditionally been concerned with reducing this lag and making the CPI more timely. From 1978 to 1994, they measured inflation using an average of one month and six month rent changes as follows:

$$I_t = 0.65r_{t-1}I_{t-1} + 0.35r_{t-6}I_{t-6} \quad (5.23)$$

This composite estimator was introduced by the BLS to achieve two goals: reduce the variance and improve timeliness. The Price Statistical Methods Division of the BLS found that tenants typically underreported 1-month changes which resulted in a bias. This recall bias was especially prevalent among new tenants who may be unaware the the rent changed when they moved in. Beginning in 1985, the BLS attempted to correct this bias by imputing 1-month changes for newly occupied units. However, even among long-term tenants, 1-month changes tended to be underestimated. In recognition of the 1-month recall bias, the BLS switched from the composite estimate to the current method using the sixth root six month changes (Armnecht, Moulton, and Stewart, 1995).

This lag in the CPI, while preferable to 1-month recall bias, nevertheless further reduces the timeliness of the CPI. The BLS argued, in changing to the current method, that the slow moving nature of rents reduces the significance of this reduced timeliness:

The only potential disadvantage is a possible reduction in the timeliness of the indexes, since it would take an average of 3-4 months for

rent changes to appear in the index.... Since rents are largely determined by long-term contracts and tend to move gradually, the disadvantage of reduced timeliness is not so great for shelter as it would be for other components of the CPI (Armnecht, Moulton, and Stewart, 1995).

However, to the extent that the slow movement of rents is due to their long-term contracts, this applies to average rents and not market rents. The underlying market rents are not affected by contract length, so the need for more timely measures is greater when measuring owners' equivalent rent than when measuring rents themselves.

5.7 Summary

Overall, Chapter 5 has shown that there are theoretical reasons to suspect that marginal rents and average rents diverge in important ways. This section showed that compared to marginal rents, average rents are biased downward and are affected by a shock with a lag. In addition, it was shown that price relative based on average rents are not biased in the long-run but are biased towards zero in the short-run. Changes in turnover have also been shown to potentially cause average rent to diverge from marginal rent. However, the literature on endogenous price changes suggests that some of these effects may be mitigated. Finally, this chapter discussed an additional lag that result from the BLS method of measuring OER.

CHAPTER 6

ALTERNATIVE MEASURE OF OER

The previous chapters have established the desirability of measuring marginal rents and the problems with average rent measures used by the CPI. This then begs the question, “How should marginal rents be estimated?”. This chapter will propose the use of a repeat-rent regression method and also introduce two datasets to test whether, and to what extent, measures of inflation using this method diverge from inflation measured using the traditional BLS CPI approach. In addition, another dataset of rentals and the repeat-rent methodology will be used to construct a measure of owner-occupied housing inflation for a particular BLS sample area. This will allow the comparison of the alternative inflation estimate with the official CPI estimate for that geography, and the use of a statistical test of inflation forecasting. In addition to being more theoretically desirable, this test will indicate if marginal rents have preferable statistical properties over average rents.

6.1 Measuring Marginal Rents

The simplest way to measure changes in marginal rents would be to use the price relative approach currently used in the CPI, except only using observations where leases and tenants in a unit are new in that month. However, generating a price relative requires there to be an observation in consecutive periods. If using the BLS panel approach, this would mean finding units that not only have a new lease in the current period but also a new lease six months ago and taking the sixth root of the price relative. Since the majority of leases are 12 months in length, this would mean finding a large dataset of tenants, throwing out most observations, and using a possibly non-random subsample.

Instead, econometric methods can be used to estimate constant quality measures

of rental inflation. The general approach to econometric models of rent inflation has been to extend approaches used to measure house price inflation to rents (SNW, 2008; Ambrose, Coulson, Yoshida, 2012). The two most common models used are hedonic regression and repeat sales.

Broadly speaking, the hedonic approach models the log of house prices as a function of housing characteristics, a vector of dummies indicating the period of sale, and an error term. The housing characteristics control for quality differences and the time dummies are used to create a price index. However, hedonic indexes have several shortcomings. One is that they can be biased if the correct functional form is not specified, or if an incomplete or incorrect set of hedonic variables is used. Furthermore, because different researchers tend to use different functional forms, housing characteristics, error specifications, and dependent variable transformations, this method is sometimes regarded as not entirely reproducible (Diewert, 2009). Collecting a complete set of housing characteristic variables can also be difficult, as some data, particularly neighborhood characteristics, is simply unavailable for many attributes that would be expected to have large impacts on the value of a house (Case, Pollokowski, Wachter, 1991).

Due to the shortcomings of hedonic models, repeat sales models are often used. This method, first described by Herman Wyngarden in 1927 and later rediscovered from obscurity by Carl Case in 1986 (Shiller, 1987), only uses information on sales that have transacted more than once during the timeperiod of the index. The primary advantage of this over hedonic models is that it forgoes specifying a relationship between housing characteristics and price. Instead, characteristics of the house are either assumed to have not changed between sales, or, when data are available, only homes that have not changed over time are used in the index. Therefore, for these homes, the changes in prices are pure appreciation that is not dependent on quality or other characteristics of the home.

Repeat sales indexes are perhaps the most common housing price indexes. The popular Case-Shiller indexes utilize a repeat sales methodology as does the Conventional Mortgage Home Price Index (CMHPI) produced by the Office of Federal Housing Enterprise Oversight (OFHEO) (Calhoun, 1996).

The repeat-rent and repeat-sales methods assume that the price of a house i at time t , P_{it} , is the product of a market price index β_t , gaussian random walk H_{it} , and white noise process η_{it} :

$$P_{it} = \beta_t H_{it} \eta_{it} \quad (6.1)$$

$$\ln(P_{it}) = \beta_t + H_{it} + \eta_{it} \quad (6.2)$$

Then the percentage change in the price of house i sold in time periods t and s , with $t > s$, is:

$$\begin{aligned} \Delta V_i &= \ln(P_{it}) - \ln(P_{is}) \\ &= \beta_t - \beta_s + H_{it} - H_{is} + \eta_{it} - \eta_{is} \end{aligned} \quad (6.3)$$

Every observation then consists of two transactions. Let $D_{i\tau}$ represent a dummy variable equal to 1 if house i transacted for a second time in time period τ , and equal to -1 if it transacted for the first time in period τ . Then equation ΔV_i can be rewritten as:

$$\Delta V_i = \sum_{\tau=0}^T (P_{i\tau}) D_{i\tau} \quad (6.4)$$

Then from the assumptions that $E[H_{it} - H_{is}] = 0 \quad \forall t, s$ and $E[\eta_{it} - \eta_{is}] = 0 \quad \forall t, s$, we can write ΔV_i as:

$$\Delta V_i = \sum_{\tau=0}^T \beta_\tau D_{i\tau} \quad (6.5)$$

This formulation allows the estimation of the repeat-rent price index I_t as:

$$I_t = e^{\beta_t} = \frac{E[R_t]}{E[R_0]} \quad (6.6)$$

6.2 Lease Level Rental Data

6.2.1 Data Description

The first dataset is comprised of 398,449 leases for multifamily units from a large Real Estate Investment Trust (REIT), which are corporations or partnerships organized for the purpose of owning, and often operating, income producing real estate. The REIT supplying the data owns and operates multi-family residential communities in the Northeast, Mid-Atlantic, Midwest, Pacific Northwest and parts of California. The dataset contains 280 multifamily communities containing 345 buildings and 82,005 units.

After data cleaning¹ 99.7% of the leases were useable. Table 6.1 below summarizes unit level characteristics that are available in the data. Importantly, the mean number

¹There are a small number of early leases, but 99.88% of them begin in 2001 or later. Pre-2001 leases are excluded from this analysis because only timeperiods with a significant number of useable observations can be used to estimate average prices. Thirty-three additional units with 268 leases are dropped because they contain more than one lease with the same start date and different ending dates. This can be a result of one tenant moving out early and another moving month-to-month, or because both residents moved out early. Because it cannot be determined which is occurring, all units with any duplicate leases are dropped. Two leases are dropped because their lease end dates are before the lease begin dates. The other leases for these units are kept because, in both cases, the erroneous lease was the first observation, and so subsequent data are assumed to be correct.

Another 2,814 units contain inconsistent square foot measures over time. The square footage listed on a lease is judged relative to the square footage on the most recent lease for that unit, as the REIT's information is expected to get more accurate over time. For 147 of these units the error is only temporarily inconsistent, meaning the square footage on the lease matches before and after the lease with mismatched data occurs. These are assumed to be coding errors and are corrected. For the remaining units, the vast majority (94%) of the changes in square foot from lease to lease correspond with a change in the data system, which strongly suggests it was a change in the coding rather than a change in the actual unit size. There are 44 units that have at least one change in square footage that does not correspond to a change in the data system. These 44 units and their 414 corresponding leases are dropped from the data.

There was also temporal inconsistency in a small number of units in terms of bedrooms and bathrooms. When the change occurred at the same time as the change in recording system, it was regarded as a coding error, and the most recent data were assumed to be correct. Otherwise, the unit and its leases were dropped. This resulted in an additional 14 units and 164 leases being dropped from the analysis.

of leases per unit is 5, meaning several changes in lease per unit can be observed.

Table 6.1: Unit Level Characteristics

Variable	Min	Max	Mean	Median
Bedrooms	0	4	1.6	2
Bathrooms	1	4	1.5	1
Square Feet	278	2,827	969	933
Building Units	1	631	89	30
Year Built	1905	2012	1994	1999
Number of Leases	1	29	4.9	4

The median unit has two bedrooms, one bathroom, 933 square feet, four leases, and is in a building built in 1999 with a total of 30 units.

One shortcoming with the data is the inability to distinguish between months where a unit is on a month-to-month lease versus when a unit is vacant, as neither are covered by the leases in the data and show up as missing information. Conversations with the REIT have indicated that vacancy is generally low, yet the data indicate that information is missing for 30% of months overall. This suggests a non-trivial proportion of the unit-months have month-to-month leases. To control for this, it is assumed that a given lease for a unit only ends when a new lease begins. Thus, for the period when information is missing, the previous lease is assumed to be in force. This increases the nominal rigidity of the series, and the data can be considered a partially simulated dataset with slightly greater nominal rigidity than in reality. However, given that the units are multifamily rental units, which have lower rigidity than single-family rentals (Verbrugge and Gallin, 2012), this is arguably a more relevant level of rigidity to consider as it has a more comparably rigid measure than a multifamily dataset with no added rigidity.

After the initial cleaning and narrowing of the sample, the data are transformed into a panel dataset of 5.4 million observations of rent by month by unit. This allows changes to be computed at various intervals and also for outlier detection to

be performed. If a unit was identified as having a percent change in rent in any month that fell outside of 45.8%², then it was removed as an outlier. This process lead to the dropping of 3,143 units, which corresponded to 284,606 unit-months of panel data observations or 5.3% of the total.

A potential complication that arises is how to treat new units that enter the sample after the beginning of the overall sample time period and units that drop out of the sample before the end. When the REIT purchased or built a unit, its leases appear in the data until the unit is demolished or sold. To ensure that differences in inflation measures are not driven by differential impact of how these units are treated, the analysis will only look at units that existed in the data during the entire 2003-2011 time period. Fortunately, this still leaves a substantial sample of units, as 17% of the units have a starting lease in 2003 or earlier and an ending lease in 2011 or later. The 13,208 units that comprise this sample are of the same order of magnitude as the CPI housing sample which consists of around 28,000 units.

The data show strong evidence of nominal rigidity. Figure 6.1 below shows the average length of a price spell or, in other words, the length of time between price changes.³

Over 17% of the price spells last exactly 12 months as is expected with annual leases. The average spell length is 15.4 months, and the median is 13. The hazard function below confirms that the percent of pricing spells that last longer than 12 months is 50%. Also plotted, in Figure 6.2, is the cumulated hazard function for single-family rentals from Gallin and Verbrugge (2012). All three exhibit similar levels of rigidity for the first year, however the the rigidity of the partially simulated rents is closer to the rigidity of the single-family rents for the lower portion of the

²This corresponds to the final threshold after two rounds of removing those outside of two standard deviations of changes. The first standard deviation was 1,130% due to a few outliers where rents went to \$1. The standard deviation was computed only among the universe of single month changes that were not equal to zero.

³The excludes the last price spell, which will suffer from censoring.

Figure 6.1: Histogram: Duration of Price Spells

distribution. This indicates that the partially simulated rents are a reasonable proxy for single-family rentals in terms of overall rigidity, and that the process of imputing missing data does not lead to an unrealistic level of rigidity.

6.2.2 Inflation Comparison

Using the 13,208 units and 1.4 million month-units of observations that have a first month before the end of 2002 and a last month in the data no earlier than the end of 2011, both repeat-rent and a proxy for the CPI method using by the BLS are estimated. Comparing these two indexes estimated using a single data series will isolate the effects of index methodology.

A simple repeat rent regression of the following form is used:

$$\ln(P_{ij}) - \ln(P_{ik}) = \sum_{\tau=0}^T \beta_{i\tau} D_{i\tau} \quad (6.7)$$

Where $D_{it} = 1$ if $j = t$, meaning it is the later lease, and $D_{it} = -1$ if $k = t$, meaning it is the earlier lease. The β_{τ} are time period coefficients used to construct price index.

Figure 6.2: Hazard Function: Percent of Units With Zero Change

The repeat rent price index I_t^{RR} is then created using the formula:

$$I_t^{RR} = e^{B_t} = \frac{E[R_t]}{E[R_T]} \quad (6.8)$$

The observations used are the beginning month of each lease rather than a panel dataset consisting of every month the lease is in effect. The first month with enough observations to estimate a price index is January, 2003, and the last is December, 2011. This produces an index equal to the ratio of the prices in any given period relative to the base period. To compare, the BLS method is used to compute price relatives from the CPI formula:

$$PR_{t,t-6} = \frac{\sum_i R_{i,s,t}}{\sum_i R_{i,s,t-6}} \quad (6.9)$$

Where $PR_{t,t-6}$ is the price relative for period t to $t-6$. The CPI price index I_t^{CPI} is estimated by moving the index I_{t-1}^{CPI} forward using the sixth root of $PR_{t,t-6}$, which

approximates the one month change:

$$I_t^{CPI} = I_{t-1}^{CPI} \cdot \sqrt[6]{PR_{p,t,t-6}} \quad (6.10)$$

In addition, to isolate the effect of the six-month lagging, a CPI proxy with no lags can be computed using the following formula:

$$PR_{t,t-1} = \frac{\sum_i R_{i,s,t}}{\sum_i R_{i,s,t-1}} \quad (6.11)$$

To estimate this index, a lease is used in all periods for which it is in effect, not just when it is new as in the repeat rent index. This means utilizing a panel dataset of the same 13,208 units, but with 1.4 million unit-month observations for every month that the lease was in effect. The difference in sizes of the datasets used in each method reflects the different approach to using the information. Repeat-rent uses the leases to impute information about the narrow period when they are new, and thus reflects market rents. In contrast, the CPI method uses the leases to impute information even in periods when the leases are old which results in each observation being used more often and provides a larger dataset. The inflation estimates that result for the two series can be seen in Figure 6.3 below.

Several results are clear from this simple comparison. First, the repeat rent index increases faster than the CPI series during the general period of the national house price bubble. Similarly, the popping of the bubble appears more quickly in the repeat rent series, and so does the recovery.

From January of 2003 to the peak in August of 2008, the repeat-rent series has increased by 22.3%. In contrast, the CPI method series with 1-month lag has increased by 13.8%, and the 6-month lag series has increased 12.3% from its first period in June, 2003. The 1-month CPI series peaks two months later, up a total of 14.4%, while the 6-month lag CPI peaks five months later at 14.4% up from the base period. By

Figure 6.3: Rent Index Method Comparison

February 2009, as deflation is just beginning in the CPI 6-month lag series, the repeat rent series has already declined 6% from the peak. The CPI 1-month lag series has declined by 1%. This suggests the CPI method reflects a bubble more slowly and underestimates the size of the appreciation that has occurred.

The bottom of the bubble occurs in the repeat-rent series in December of 2009 while the bottom occurs four months later in the CPI 1-month lag series, and 6 months later in the CPI 6-month series. From peak to trough the repeat-rent measure has taken 15 months, and prices have fallen 11.3%. In the CPI 1-month lag series, peak-to-trough lasts 18 months and prices have fallen 7.8%. For the CPI 6-month lag series, peak-to-trough is 17 months with a decline of 7.9%. Owing in part to the larger decline in prices, the recovery is faster and stronger in the repeat-rent series. From trough to the end of the data series in November, 2012, the repeat-rent series increases 13% while the CPI 1-month and CPI 6-month increase 8% and 7% respectively. Finally, in the last six months, the repeat-rent series has begun to decline again, suggesting

future declines which have not yet shown up in either CPI measure.

Similar results can be seen by looking at 12-month percent changes, shown in Figure 6.4 below for the repeat-rent index and the CPI 6-month lag. The growth rates begin to decline in September 2006 for the repeat-rent series after peaking at just under 8%. For the CPI series, this deacceleration does not begin until 26 months later, in November 2008, after peaking at 5.3% the month before. The month after the CPI series first begins deacceleration, the repeat-rent series is showing deflation with a rate of -1%. Deflation does not appear in the CPI series until 9 months later, in August 2009. In this same month that the CPI series is first showing deflation, the repeat-rent growth rate hits its bottom at nearly -10%. The CPI series will not bottom until 8 months later in April 2010, in the same month that the repeat-rent series deflation ends.

Overall, the repeat-rent series shows deflation and inflation significantly sooner than the CPI method, either with the 6-month lag currently used by the BLS or with a 1-month lag. In addition, the rise and fall of prices is more stark in the repeat-rent measure. The repeat-rent series is more volatile, but within the same order of magnitude. The standard deviation of 12-month changes is 0.047 for the repeat-rent method, and 0.034 for the CPI method. These results suggest that there are significant gains in timeliness to be had from using a marginal rather than average rent measure for OER.

6.3 MLS Rental Data

6.3.1 Data Description

The use of REIT lease data allows for the comparison of the CPI method with a repeat-rent method using the same data, which has the advantage of isolating the effects of the index methods from effects that might result from different data sources. However, the data did not contain geographic information, so it is not known how

Figure 6.4: Rent Index Method Comparison, 12-Month Percent Change

representative of the U.S. as a whole the data is, nor how representative of any particular geographic subregion. As a result, the estimated marginal rent inflation series is not useful for comparing to the actual CPI as calculated by the BLS. A more representative dataset would also be useful in order to determine whether the marginal rent series outperformed the actual CPI in terms of forecasting overall inflation. In order to test this, a new dataset of rents that are geographically concentrated in one BLS sampling area will be used to compute a marginal rent series to compare to the official BLS CPI for that area.

This exercise will show whether the marginal rent series reflects inflation faster than the official CPI measurement as predicted by the results in the previous section. In addition, it will offer several advantages over the comparison of marginal rent and CPI proxy using lease data in the previous section.

First, the use of the official CPI series rather than a proxy for comparison means all of the official adjustment procedures that were foregone in the previous section

will be utilized. This includes vacancy and utilities adjustments. If these procedures speed up the transmission of underlying shocks, then the official CPI rent inflation may more closely resemble the marginal rent inflation than the proxy CPI method measure does.

The second advantage of this dataset is that it is based on rents for single-family housing units, which makes it more similar to the owned housing stock for which housing inflation is being proxied.

In addition, this data set also allows for an important extension in allowing for a statistical test comparing marginal rents to the actual CPI as an inflation measure. Given the Federal Reserve's need to be forward looking, previous research has compared inflation measures by their ability to forecast future inflation (Crone, Khettry, Mester, and Novak, hereafter CKMN, 2011). The estimated marginal rent inflation series can be compared to official inflation estimates to see if the new measure better forecasts overall future inflation. In addition to the theoretical arguments for using marginal rent discussed in previous chapters, if marginal rents are found to better forecast overall inflation, it will be evidence in favor of the BLS adopting this measure.

The dataset for this comes from MRIS, a large Multiple Listing Service (MLS) in the Mid-Atlantic area. MLSs are organizations usually run by one or more local boards of Realtors for the purposes of providing a web portal Realtors can use to list and search for housing for sale. Datasets compiled from MLS transactions have been a rich source of information for real estate economics studies. A recent meta-analysis of 71 housing price studies found that 45% of the studies utilized MLS data (Sirmans, MacDonald, Macpherson, and Zietz, 2006). Some MLSs also collect and make available on their web portals data on rental units, because Realtors will sometimes list rental units for a landlord client. This is the second study to date to utilize MLS rental data.⁴

⁴One other example of the MLS rental data used by economists is a short note published by Core Logic using their proprietary MLS rental dataset which has not yet been made publicly available.

The MRIS dataset has several quality advantages over other listing data. Importantly, Realtors utilizing the MLS are required to enter a final lease price in addition to a listing price. In addition, because MRIS uses the rental data for their own analysis and marketing reports, they perform quality control measures and contact Realtors who have posted a listing if the data appear suspicious.

The dataset contains 276,158 single family unit listings from 2000 to 2012, 41.1% of which are in Maryland, 54.7% are in Virginia, 2.6% are in Washington D.C., with the small remainder scattered throughout other Mid-Atlantic states.⁵ Each listing contains the following information on the unit: address, building type, number of bedrooms, number of bathrooms, year built, and total unit square footage. Each listing also contains the following information on the unit's lease: date listed, original listing price, final leased date, final leased price, lease length, required security deposit amount. The dataset also contains a string variable that indicates what services are included in the price of rent that can be parsed to produce dummy variables indicating if a unit's rent includes various amenities, such as heating or parking.

Due to the data's significant coverage of a specific geographic region, a repeat-rent index can be estimated that can be compared to the official BLS estimate of owner-occupied housing inflation for that same area. As discussed in Chapter 2, the BLS surveys 87 geographic areas known as Primary Sampling Units (PSUs) for its housing survey. The PSU A312 consists of 21 counties in DC, WV, MD, and VA. As Table 6.2 below indicates, there are 128,106 repeat-rent observations in the 21 counties in PSU

See Khater (2012).

⁵There are 90,132 that are for multi-family or other non-single family unit types that are not included in the analysis. These are excluded because different units in the same address cannot be tracked over time. This provides the added benefit of more closely matching the predominantly single-family universe of owner-occupied housing for which OER is being imputed. In addition, there were a small percentage dropped for data cleaning purposes. Dropped observations include: 537 with lease terms over 10 years, 99 with unit size over 10,000 sf, 13 with missing rent or list price data, 788 where rental price was less than 50% or greater than 150% of the list price, and a single unit with a list price of \$1. Finally, when there was more than one observation for a unit in a given month, the last listing was kept, which resulted in 552 duplicate observations being dropped. Overall, 1,991 or 0.7% of the single-family units were dropped due being outliers or duplicates.

A312. This coverage allows for the creation of a repeat-rent estimate of marginal rent inflation for PSU A312.

Table 6.2: Repeat-Rent Observations in PSU A312

County	State	Repeat-Rent Observations	Total OOH Units	County Share of Repeat-Rent	County Share of OOH
District of Columbia	DC	3,185	111,879	2.5%	8.2%
Calvert	MD	2,491	25,754	1.9%	1.9%
Charles	MD	3,966	40,454	3.1%	2.9%
Frederick	MD	2,333	64,112	1.8%	4.7%
Montgomery	MD	17,907	244,815	14.0%	17.8%
Prince George's	MD	8,065	194,047	6.3%	14.1%
Washington	MD	544	36,741	0.4%	2.7%
Arlington	VA	4,637	43,168	3.6%	3.1%
Clarke	VA	178	4,195	0.1%	0.3%
Culpeper	VA	846	11,498	0.7%	0.8%
Fairfax	VA	40,221	274,448	31.4%	20.0%
Fauquier	VA	2,109	17,633	1.6%	1.3%
King George	VA	484	6,261	0.4%	0.5%
Loudoun	VA	11,787	77,022	9.2%	5.6%
Manassas	VA	896	8,003	0.7%	0.6%
Prince William	VA	15,264	93,372	11.9%	6.8%
Spotsylvania	VA	4,654	32,759	3.6%	2.4%
Stafford	VA	6,626	31,502	5.2%	2.3%
Warren	VA	372	10,591	0.3%	0.8%
Berkeley	WV	923	29,297	0.7%	2.1%
Jefferson	WV	618	14,903	0.5%	1.1%
Total		128,106	1,372,454	100%	100%

Compared to typical repeat-sales datasets, there are a large number of repeat-rent observations in the MRIS data. The repeat-sales studies summarized in Table 6.4 list repeat-sales percentages that make up 3.2% to 14.5% of the data. As the table below shows, repeat-rent observations make up over 60% of the cleaned MRIS data, illustrating that most observations can be used in a repeat-rent analysis. Only 76,337 out of the total of 204,443 cleaned transactions are for individual units that have no repeat-rents which leaves the 128,106 usable observations.

Table 6.3: Count of Transactions by Number of Repeat-Rents

Num. Per Unit	Count of Obs.	% of Obs.	Cumulative %
0	76,337	37.34	37.34
1	53,664	26.25	63.59
2	34,110	16.68	80.27
3	20,456	10.01	90.28
4	11,010	5.39	95.66
5	5,376	2.63	98.29
6	2,275	1.11	99.41
7	816	0.40	99.80
8	288	0.14	99.95
9	100	0.05	99.99
10	11	0.01	100.00
Total	204,443	100	

6.3.2 Inflation Comparison

The basic repeat-rent estimate from section 6.1 can be used to estimate a price index for January, 2000 through September, 2012. Figure 6.5 below displays the estimated index alongside the CPI for owners' equivalent rent of primary residence in Washington-Baltimore, DC-MD-VA-WV CMSA.

The most obvious difference between the two series is the implausibly fast growth in 2001, at times exceeding 10%. This fast growth appears in both a hedonically estimated index and a simple median rent index, so it does not appear to be either a quality adjustment or repeat-rent methodological issue.

Further visual inspection of the two indexes in Figure 6.5 reveals several things. First, while both series are not seasonally adjusted, the repeat-rent series displays higher seasonality. The series diverge at first, with the repeat-rent series growing faster, then converge by 2007 when the CPI grows more quickly, and finally converge again by 2011.

Looking at the series in one-month changes, the differences between repeat-rent and the CPI series become more stark. The repeat-rent series is far more volatile, with a standard deviation of .0082 compared to .0026 for the CPI series. Given that

Figure 6.5: Repeat-Rent Index Compared to CPI for OER

the CPI series uses a six-month averaging, its lower volatility is unsurprising. Figure 6.7 below illustrates the effect of a six-month smoothing on the repeat-rent series. It is closer in volatility to the CPI, with a standard deviation of .0044. Even with the smoothing, this is still nearly double the standard deviation of the CPI series. The increased volatility of the repeat-rent index and changes in the index is consistent with the theory in Chapter 5.

To abstract from the issue of seasonality in each series, year-on-year changes, a common form in which CPI changes are reported, can be examined. As seen in Figure 6.8 below, the difference in volatility is only slight, with the CPI having a standard deviation of 0.0127, and the repeat-rent series having a standard deviation of 0.0198. In fact, if the first year of idiosyncratic high growth in the repeat-rent is excluded by looking at the 2002 and forward standard deviations, the repeat-rent series is less volatile than the CPI, at 0.0127 versus 0.0128 respectively.

The year-on-year changes illustrate several crucial trends for the DC/Baltimore

Figure 6.6: Repeat-Rent Compared to CPI for OER
One-Month Changes

area. First, as indicated by the levels and contrary to the prediction of theory, the housing bubble did not show up more starkly in the repeat-rent series. In fact, while the CPI increased from an average of 4% inflation to 6% inflation from 2006 through 2007, the repeat-rent series remained steady. However, between 2003 and 2006, the repeat-rent series did have a gradually increasing inflation rate, which nevertheless remained below or near the CPI inflation over this period.

Another crucial pattern illustrated in the year-on-year graph is that the repeat-rent series reflects market turning points more quickly. These can be seen by comparing the two series to the Case-Shiller house price index for Washington DC, which is shown in Figure 6.9 below.⁶

The axis on the right provides the scale for the HPI, and the axis on the left for the CPI and repeat-rent measures. While Case-Shiller tracks closely to the repeat-rent

⁶To extrapolate from the idiosyncratic early growth in the repeat-rent series, only 2003 and forward are shown.

Figure 6.7: Smoothed Repeat-Rent Compared to CPI for OER
One-Month Changes

measure, the variation is much larger in the HPI.

The Case-Shiller index begins the deacceleration that marked the start of the end of the housing bubble in May 2005 as it declines from the peak growth rate of 26.8% in the previous month. The CPI begins deacceleration one year and seven months later in December 2006, a month after achieving the peak bubble period year-on-year growth rate of 6.9%. The repeat-rent index begins deaccelerating two months earlier than the CPI, beginning the decline in October from the previous month's peak of 4.9%.

The bottom of the housing bust is reflected even more quickly in the repeat-rent series. The second derivative of the HPI changes in January 2009 as prices start declining at a slower pace. The previous month was the largest decline in the housing bust, at -19.6%, and from January on, things began declining at a slower rate. The CPI did not stop deaccelerating until July, 2010. The previous month, it reached the low point of its growth rate at 0.4%, and began growing increasingly quickly after that

Figure 6.8: Repeat-Rent Compared to CPI for OER
Twelve-Month Changes

point. The repeat-rent series began turning around in May of 2009, fourteen months before the CPI. It reached its lowest growth rate the previous month of -0.6%, and conditions began improving thereafter.

The recovery stall in late 2010 shows up in the repeat-rent series very close to when it shows up in the HPI, but is not reflected in the CPI until around a year later.

6.4 Repeat-Rent Shortcomings and Extensions

One common criticism of the repeat-sales approach is that by only using observations that transact more than once in the dataset, a large amount of information is thrown out. The following table from Nagaraja, Brown, and Wachter (2010) shows the amount of data that is lost by using repeat-sales in two studies and a handful of cities.

However, given the much lower average tenure length for rental units than for owner-occupied housing units, the proportion of units that transact more than once

Figure 6.9: Repeat-Rent Compared to CPI for OER and Case-Shiller for D.C., Year-on-Year Changes

Note: CPI and Repeat-Rent estimates use the left axis, and the Case-Shiller for D.C. uses the right axis.

is likely to be much higher in a dataset of rentals. If the same tenant with a new lease is considered a repeat observation, then the majority of units should have repeat observations in a dataset with even a few years of data. The 2011 American Housing Survey reported that as of 2011, 41% of renters had only lived in their homes since 2010 compared to 8% for owners. In addition, 82% of renters moved into their homes within the last 6 years compared to 33% for owners. Therefore, regardless of whether new leases or new units are considered repeat-rent observations, the loss of data should be significantly less of a problem than for repeat-sales.

Another potential issue with repeat sales estimates in the housing price index literature is heteroskedasticity. Recall that the dependent variable of the repeat-rent regression is the log difference of lease price for a particular unit i sold in time periods

Table 6.4: Effect of Removing Single Sales on Sample Size

City	No. Observations	No. Repeat Sales	% Repeat Sales
Case-Shiller (1970-1986)			
Atlanta, GA	221,876	8,945	4.0%
Chicago, IL	397,183	15,530	3.9%
Dallas, TX	211,638	6,669	3.2%
San Francisco/Oakland, CA	121,909	8,066	6.6%
Meese-Wallace (1970-1986)			
Freemont, CA	23,408	3,405	14.5%
Oakland, CA	27,606	3,342	12%

Source: Nagaraja, Brown, and Wachter (2010)

t and s , with $t > s$, is:

$$\begin{aligned}
 \Delta V_i &= \ln(P_{it}) - \ln(P_{is}) \\
 &= \beta_t - \beta_s + H_{it} - H_{is} + \eta_{it} - \eta_{is}
 \end{aligned} \tag{6.12}$$

Case and Shiller (1987) pointed out that heteroskedasticity can occur if the assumption fails that $E[(H_{it} - H_{is})^2] = \sigma^2 \quad \forall t, s$. Instead, they argue it is likely the error variance is a function of the time between the two observations:

$$E[(H_{it} - H_{is})^2] = \rho(t - s) + \phi(t - s)^2 + \mu \quad \forall t, s, i \tag{6.13}$$

To correct for this they propose a 3-step GLS procedure. First, a basic repeat-rent regression is estimated. Then, for each observation ΔV_i with a first lease in period s and second lease in period t , the residual ϵ_{its} is estimated. Then a second regression is performed using ϵ_{its}^2 as the dependent variable and the time between the two observations ($t - s$) as the independent variable, both linearly and quadratically. The predicted values from this regression are then used as weights, and the repeat-rent is run again as weighted least squares.

While the GLS model is more efficient, as Figure 6.10 below shows, it does not substantially alter the results for the MLS data. Twelve month inflation rates are largely similar, with the GLS inflation rate average at 3.2%, and the basic repeat-rent measure averaging 3.1%. Volatility is actually increased as a result of the GLS correction, going from a standard deviation of 0.018 to 0.019.

Given the lack of improvement in volatility and small change in average year-on-year growth, the added complexity of the GLS correction does not appear to be justified.

Figure 6.10: Basic and GLS Repeat-Rent Comparison
Year-on-Year Percent Changes

Another possible extension is to control for the fact that the CPI is an arithmetic mean, and the repeat-rent estimation is geometric mean. Shiller (1991) provides a method to estimate an arithmetically weighted repeat-sales, however Goetzmann (1992) proposes a simple adjustment that can be made to the GLS repeat-sales approach that approximates the arithmetic results. Using the coefficients $\tilde{\beta}_t$ from the GLS repeat-rent, the arithmetic index is:

$$I_t^A = e^{\hat{\beta}_t + \frac{1}{2}\hat{\sigma}_t^2} \quad (6.14)$$

Where $\hat{\sigma}_t^2$ is estimated using the second stage GLS coefficients:

$$\hat{\sigma}_t^2 = \hat{\rho}(t - t_0) + \hat{\phi}(t - t_0)^2 \quad (6.15)$$

Therefore, in the baseline period, t_0 , no adjustment is made. In each subsequent period, the adjustment is based on elapsed time relative to the baseline. However, using the MLS data the estimated coefficients for equation 6.15 are small, and the adjustments only trivially impact the index. In the last period of the index, when the difference between arithmetic and standard indexes should be greatest, the difference is approximately 0.01%.

In this data sample, neither the GLS correction nor the arithmetic weighting correction alter the index in any meaningful way. Given a goal of the BLS is to make the CPI as clear and comprehensible as possible, these results suggest an unadjusted repeat-rent regression is optimal.

6.5 Policy Implications

6.5.1 Overview

The previous sections used two datasets to document that repeat-rent estimates of rent were more timely than estimates using the CPI method; or in other words, marginal rents are more timely than average rents. However, it is unclear what the policy implications of this timeliness are. This section will provide two pieces of evidence that suggest policy improved if the BLS moved to a marginal rent measure. First, using the MLS rental data, econometric evidence will be used to show that the marginal rent measure is superior to average rent for forecasting purposes. In addition, using both datasets, the implications of switching to marginal rents for

monetary policy will be shown by examining what the Taylor Rule implies under marginal versus average rents. This evidence suggests that monetary policy could have been more responsive to the housing bubble and recession if marginal rents had been used instead of average rents.

6.5.2 *Forecasting Comparison*

While visual inspection provides evidence that the repeat-rent series is able to reflect underlying market changes in a more timely manner, econometric evidence can also be brought to bear. If it can be shown that a particular measure of inflation better predicts total future inflation than total future inflation predicts itself, then this is evidence that the measure is more useful for necessarily forward looking monetary policy. Past research has focused on comparing different measures of core inflation to overall inflation. Blinder and Reis (2005) look at core CPI's ability to forecast overall CPI, whereas CKMN include the additional measures of CPI less energy, the Cleveland Fed's weighted median CPI, and also similarly analyze PCE inflation. In light of the past literature, there are two potential analysis that can be done to determine whether repeat-rent provides a superior measure of OER.

First, it will be determined whether repeat-rent forecasts current OER better than current OER forecasts itself. Additionally, it will be seen which measure better forecasts overall inflation. It could be, for instance, that repeat-rent is a better predictor of the current OER measure, but that the current OER measure better forecasts overall inflation due to the added volatility of repeat-rent.

Several studies support the use of univariate inflation forecasts (CKMN; Ang, Bekaert, and Wei, 2007; Stock and Watson, 2007; and Blinder and Reis, 2005) which greatly simplifies the exercise compared to a multivariate model. Following CKMN and Blinder and Reis (2005), the following univariate model will be employed:

$$\pi_{t,t+h} = \alpha + \beta X_{t-12,t} + \epsilon_t \tag{6.16}$$

Where $\pi_{t,t+h}$ is the percentage change in total inflation from t to $t+h$ for PSU A312, and $X_{t-12,t}$ is the twelve month percentage change from $t-12$ to period t in either the marginal rent inflation series or the official BLS CPI for owners' equivalent rent for PSU A312.

Following Reis and Blinder (2005) and CKMN, future inflation will be forecast at intervals of $h = 6$ months, 12 months, 24 months, and 36 months. Rolling regression forecasts will be computed using a window of 60 months for estimation. For example, the earliest forecast period is January 2006. The model for this month was estimated using the 60 months from January 2001 through December 2005. For the next forecasted period, February 2006, the model was reestimated using February 2001 through February 2006. A rolling forecast as used by CKMN has two advantages over the fixed period forecast used in Reis and Blinder (2005): it allows parameters to change over time reflecting underlying structural change, and it reduces the influence of parameter estimation noise (CKNM).

Three measures of overall inflation are used for $\pi_{t,t+h}$. First, owners' equivalent rent of primary residence in PSU A312 is used, which is also the second independent variable. This is used to determine whether marginal rent can forecast OER better than it can forecast itself. In addition, two measures of overall inflation are used: overall CPI for the U.S. and overall CPI for the Southern Census region, which includes PSU A312.⁷

Figures 6.11 through 6.13 show the forecasts and the actual values for the three measures of inflation and four forecast horizons. Overall, visual inspection suggests that forecasts using repeat-rent tend to more closely track actual values than forecasts using OER.

⁷Inflation for the PSU A312 is the most applicable measure of overall inflation, however for this area overall inflation is only measured every other month.

Figure 6.11: Forecast Comparison
 $\pi_{t,t+h} = \text{OER}$, Washington-Baltimore CMSA

Figure 6.12: Forecast Comparison
 $\pi_{t,t+h}$ = Overall CPI, National

Figure 6.13: Forecast Comparison
 $\pi_{t,t+h}$ = Overall CPI, South Region

A first test of forecasting accuracy is to compute root mean squared errors for each forecast. In addition, to test for differences in forecasting accuracy between the two inflation measures, the Giacomini-White statistic for differences in mean squared errors will be used. This statistic compares a baseline forecasting model with an alternative using the following statistic:

$$\lim_{n \rightarrow \infty} \frac{\frac{1}{n} \sum (\varepsilon_{b,t}^2 - \varepsilon_{a,t}^2)}{\sqrt{\sigma^2/n}} \xrightarrow{d} N(0, 1) \quad (6.17)$$

Where $\varepsilon_{b,t}$ is the residuals from the baseline model, and $\varepsilon_{a,t}$ is the residual for the alternative model. In this exercise, the CPI measure of OER will be the baseline, and the repeat-rent measure will be the alternative.

Following CKMN, the Newey-West method to correct for autocorrelation is used to estimate the variance using a lag of $h - 1$. The result is a two-sided test statistic. If the baseline model forecasts better than the alternative, then the squared errors for the baseline will be smaller than the alternative on average, and the test statistic will be negative. If the alternative outperforms the baseline, then the opposite will be the case. Therefore, positive and significant value for the GW statistic suggests the repeat-rent measure performs better, and a negative and statistically significant value suggests the CPI measure performs better.

As shown in Table 6.5 below, the repeat-rent measure of inflation generally produces more accurate forecasts using RMSE criteria and GW tests. The RMSE for the repeat-rent measure is smaller in all but one case, when the two are equal. In ten out of twelve GW tests, the repeat-rent measure statistically significantly outperforms the CPI measure, and in the other two cases the GW statistic is positive, indicating repeat-rent performs better, but is not statistically significant. Importantly, the CPI measure never performs better than repeat-rent.

Overall, the evidence suggests that the repeat-rent measure performs better at forecasting total inflation.

Table 6.5: Measures of Forecasting Performance

Forecasting OER							
h	RMSE: OER	RMSE: RR	Mean ε^2	Diff.	GW	p-value	N
6	1.934	1.683		0.911	6.368	0.000	75
12	1.432	1.207		0.593	4.287	0.000	69
24	1.245	0.901		0.738	9.896	0.000	57
36	1.322	1.022		0.703	9.457	0.000	45

Forecasting Overall CPI, National							
h	RMSE: OER	RMSE: RR	Mean ε^2	Diff.	GW	p-value	N
6	3.376	3.339		0.250	3.206	0.001	75
12	1.852	1.744		0.388	3.827	0.000	69
24	1.300	1.133		0.407	9.657	0.000	57
36	0.932	0.932		0.000	0.018	0.493	45

Forecasting Overall CPI, South Census Area							
h	RMSE: OER	RMSE: RR	Mean ε^2	Diff.	GW	p-value	N
6	3.651	3.622		0.208	2.339	0.011	75
12	2.030	1.918		0.445	3.405	0.001	69
24	1.400	1.201		0.518	11.116	0.000	57
36	0.948	0.944		0.007	0.382	0.352	45

6.5.3 Taylor Rule

One way to gauge the importance of the proposed methodological change is to examine what the potential policy implications would be. The CPI is considered an important measure of inflation that is watched by the Federal Reserve, and the Taylor Rule is often considered useful for understanding the relationship between inflation and Fed policy: both as a description to approximate how the Federal Reserve behaves, and as a prescription for how they should behave to conduct optimal monetary policy (Asso, Kahn, Leeson, 2007). Therefore, one possible measure of policy implications would be to look at what the Taylor Rule implies Federal Reserve policy would have or should have been if marginal rent had been used instead of the existing average rent CPI measure.⁸ To do this, a measure of overall inflation for

⁸Since 2000, the FOMC has emphasized core PCE rather than CPI in its deliberations. However, the exercise is still informative as OER in the PCE is based on the same housing sample as the CPI. In addition, Taylor (2007, 2009) also uses CPI in the Taylor Rule.

the D.C./Baltimore CMSA can be computed using the repeat-rent measure of OER. Then the federal funds rate implied by this measure of inflation under the Taylor Rule can be compared to the rate implied under the actual CPI measure. This will provide an illustrative example of the change in policy that might result if marginal rent measures of OER were adopted.

Ideally, a national measure of overall inflation would be estimated. However, the marginal rent measure estimated above is only for the D.C./Baltimore CMSA. Therefore, this requires the assumption that the national level CPI would show similar relative changes in inflation from adopting a marginal rent measure as the D.C./Baltimore CMSA shows. To estimate overall inflation with marginal rent for OER, the repeat rent index above is combined with the overall CPI less shelter for the D.C./Baltimore CMSA and rent of primary residence for the D.C./Baltimore CMSA.⁹ This is estimated as:

$$CPI_t^M = \omega_\rho \rho_t + \omega_{rr} rr_t + \omega_c C_t \quad (6.18)$$

Where ω_{js} are the weights corresponding to rent of primary residence ρ_t , repeat-rent estimate of OER rr_t , and overall CPI less shelter C_t . The weights used are the average of the 2002 through 2010 weights used in the actual CPI.¹⁰ Figure 6.14 below shows the average bi-monthly changes in inflation from the past twelve months using both series. While the actual CPI is slightly higher in the beginning, the series generally track closely until the housing bubble bursts in 2006, and the repeat-rent based estimate falls sooner and more quickly. In 2009 the repeat-rent series reflects deflation four months before the actual CPI and has a lower trough. Then throughout 2010 the repeat-rent series shows higher inflation until 2011 when they again converge.

These differences in inflation can be used to estimate differences in the federal

⁹Rent of primary residence is the other main component of shelter.

¹⁰In 2009 and 2010 the weights for OER include OER for non-primary residences which are around 1.4 percentage points of the total 29 percent weights.

Figure 6.14: Overall CPI Comparison
12-Month Average Changes, D.C./Baltimore CMSA

funds rate implied by the Taylor Rule. The basic Taylor rule from Taylor (1993) is:

$$F = \zeta + \pi + .5y + .5(\pi - \bar{\pi}) \quad (6.19)$$

Where F is the federal funds rate, ζ is the real interest rate, π is the 12-month average inflation rate, y is the percentage output gap, and $\bar{\pi}$ is the long-run inflation target. Rather than estimating the actual level of federal funds, the relevant estimate is the difference in federal funds implied by the difference in measured inflation. Therefore, the following can be used to abstract from issues of selecting an output gap or long-run inflation target:

$$\frac{dF}{d\pi} = 1.5 \quad (6.20)$$

$$dF = 1.5 d\pi \quad (6.21)$$

$$dF = 1.5 (\pi_{RR} - \pi_{CPI}) \quad (6.22)$$

Figure 6.15 below shows the change in the Taylor Rule federal funds rate that would result if the CPI used a marginal rent instead of average rent measure. Initially, the new measure prescribes a lower federal funds rate, but from 2003 to 2006, the change implied gradually increases until just prior to the housing bubble collapse of 2006, when the federal funds rate should be 1 percentage point higher. This suggests that the fed should have risen rates in late 2005 to slow the housing bubble.

Then by March of 2006 this measure of the CPI implies that the federal funds rate should have begun declining. In contrast, the actual federal funds rate at this point was steady and did not begin declining until June 2007. Therefore, the use of the marginal rent measure would have meant that the monetary policy response to the popping of the housing bubble and recession occurred 15 months earlier. This would have led to monetary policy preceding the onset of the recession by 21 months instead of 6 months.

While the marginal rent measure implies the federal funds rate should have been higher starting November 2009, the actual federal funds rate was at a lower bound at this point. Therefore it is unclear whether the higher rate prescribed by the marginal rent measure and the Taylor Rule would actually have gone above the lower bound.

The data used in this exercise may be representative of the DC/Baltimore area but it is not nationally representative. In contrast, the rental data used in sections 5.3 and 5.4 is not likely to be representative of any particular U.S. geography or the single-family rental market, but is more likely to be geographically representative of the U.S. overall. Therefore, it is useful to compare what the marginal and average rents estimated there imply for the Taylor Rule. To focus on the difference in the prescribed federal funds rate would result from adopting marginal rents the following equation derived from the Taylor Rule:

$$dF = 1.5 (\pi_{RR} - \pi_{CPI}) \tag{6.23}$$

Figure 6.15: Percentage Point Change in Taylor Rule Implied Federal Funds Rate From Adopting Marginal Rent, Using MLS Data

Where π_{RR} and π_{AR} are the 12-month changes in overall CPI based on repeat-rent and average-rent measures of OER, respectively. These are estimated as:

$$\begin{aligned}\pi_{RR} &= W_{OER}\pi_{CLH} + (1 - W_{OER})RR \\ \pi_{AR} &= W_{OER}\pi_{CLH} + (1 - W_{OER})AR\end{aligned}\tag{6.24}$$

Where W_{OER} is the expenditure weight of OER in the CPI. Here the expenditure weight is approximated with a constant 30%, and so equation 6.23 showing the difference in Taylor Rule implied federal fund rates due to different OER measures simplifies to:

$$dF = 1.5 (\pi_{RR} - \pi_{CPI}) \quad (6.25)$$

$$dF = 1.5 ((0.7\pi_{CLH} + 0.3RR) - (0.7\pi_{CLH} + 0.3AR)) \quad (6.26)$$

$$dF = 0.45 (RR - AR) \quad (6.27)$$

The results, shown in Figure 6.16 below, suggest that the federal funds rate would have been higher from 2004 through October 2007. This again suggests room for the fed to counteract the housing bubble by raising rates prior to its popping, with the federal funds rate as much as 2

The gap between the repeat-rent implied federal funds rate and the CPI funds rate begins closing in the run-up to the recession. Contrary to the MLS data results, the federal funds rate here is higher in the months leading up to the recession than it otherwise would have been. However, in this same period the federal funds rate falls from a higher level down to about where it otherwise would have been. Therefore, while the use of repeat-rent implies the federal funds is slightly higher in the months up to the recession, it is falling faster and is lower at the onset.

The repeat-rent measure implies a lower rate until January, 2010 when it again implies a higher rate than the average-rent measure. However, the actual federal funds rate reached an effective lower-bound by December 2008, at which point the lower-bound means there is no effective difference between the two. Therefore, while the repeat-rent measure prescribes a relatively higher federal funds rate after January 2010, it is unclear whether this would bring the rate above the effective lower-bound.

Overall, given that the recession began in December of 2007 these results imply that the use of repeat-rent would have led to a federal funds that was higher during the housing bubble, and was slightly higher in the months prior to the recession but was falling faster. The federal funds rate would then be lower at onset and beginning

of the recession and throughout 2008 until the lower bound was in effect. After December 2008, the effective lower bound means it is unclear whether the use of either OER measure would have had any difference.

Figure 6.16: Percentage Point Change in Taylor Rule Implied Federal Funds Rate From Adopting Marginal Rent, Using REIT Data

In summary, the Taylor Rule combined with the repeat-rent measure from the D.C./Baltimore MLS data suggests that using repeat-rent CPI may have led to the Federal Reserve being more proactive in response to the recession. Alternatively, as suggested by the REIT data, the federal funds rate may have been higher in the months prior to beginning of the recession but declined faster and been lower throughout the onset of the recession and through 2008. Overall, the results suggest the use of a repeat-rent measure of OER would lead to an a more proactive response to the recession.

In addition, both data sources suggest some room for the Federal Reserve to have raised rates prior to the recession during the housing bubble period. While the MLS

data counterintuitively suggests lower rates for the 2003 through 2006 period, it does suggest rates should have been 1 percentage point higher just prior to the housing bubble collapse. In addition, the REIT data suggests higher for the 2004 through 2006 period, with the change above one percentage point most of the time and reaching as high as two percentage points in August of 2006.

Both exercises also suggest that, in the period after the recovery, the optimal federal funds rate would have been higher if repeat-rent were used. However it is unclear whether this difference would have been enough to increase the rate above the lower bound.

Caution in interpreting these results is in order. The first exercise used the MLS data which is not geographically representative of the entire U.S. but reflects conditions in the D.C./Baltimore metro area. The second exercise uses REIT data which is more representative of the overall U.S. but applies to the multifamily rental market and not the more relevant single-family rental market. However, both exercises suggest that the more timely repeat-rent measure of owner-occupied housing inflation may improve the responsiveness of monetary policy. Overall these results should be interpreted as illustrating the significance of the choice of marginal versus average rent for OER for inflation measurement, and not strong evidence that Federal Reserve should have behaved different. This is an issue that researchers should investigate further.

6.6 Conclusions

This chapter has concluded the analysis of marginal rent by demonstrating several empirical advantages. The REIT data illustrates the potential for marginal rent measures to show housing bubbles and busts more quickly and more starkly. In contrast, the comparison of the marginal rent measure and CPI measure of OER for the Washington D.C./Baltimore area suggests that the advantage of marginal rents

is primarily one of timing rather than the overall size of the price changes in the bubble and bust. While the evidence is mixed as to whether marginal rent would have allowed the national housing market bubble to have been reflected in larger changes in OER inflation, it provides clear evidence that it would have more closely tracked the timing of the bubble.

While the strongest arguments in favor of marginal rent are based on the theoretical justification as a cost-of-living index, it remains an important measurement of the price level for policymakers, individuals, and businesses. As such, the extent to which marginal rent forecasts future inflation better than current methods is an additional argument in its favor. This section provides visual evidence that marginal rent is a more timely measure, and statistical evidence of superior forecasting performance. Overall, the empirical evidence in favor of a marginal rent measure of OER based on repeat-rent supports the theoretical conclusions. This suggests the BLS should investigate the cost and practicality of implementing a marginal rent measure of OER.

Counterfactual exercises using the Taylor Rule provide suggestive evidence that if the BLS were to make this change, the implications for policymakers could be consequential. Inflation measures produced using the MLS and REIT data both suggest that marginal rent measures would lead the Federal Reserve to be more proactive both prior to and in the aftermath of the housing bubble.

BIBLIOGRAPHY

- Ambrose, B. W., Coulson, N. E., & Yoshida, J. (2012): *The Repeat Rent Index* (Working Paper).
- Armknrecht, P. A., Moulton, B. R., & Stewart, K. J. (1995). *Improvements to the food at home, shelter, and prescription drug indexes in the US consumer price index*. US Department of Labor, Bureau of Labor Statistics, Office of Prices and Living Conditions.
- Ang, A., Bekaert, G., & Wei, M. (2007). Do macro variables, asset markets, or surveys forecast inflation better?. *Journal of Monetary Economics*, 54(4), 1163-1212.
- Asso, P. F., Kahn, G. A., & Leeson, R. (2007). *The Taylor rule and the transformation of monetary policy* (No. RWP 07-11). Federal Reserve Bank of Kansas City.
- Bailey, M. J., Muth, R. F., & Nourse, H. O. (1963). A regression method for real estate price index construction. *Journal of the American Statistical Association*, 58(304), 933-942.
- Barker, D. (2003). Length of residence discounts, turnover, and demand elasticity. Should long-term tenants pay less than new tenants? *Journal of Housing Economics*, 12(1), 1-11.
- Blinder, A. S., & Reis, R. (2005). Understanding the Greenspan standard. *The Greenspan Era: Lessons for the Future*, 25-27.
- Boskin, M. J. (1996). *Toward a more accurate measure of the cost of living: final report to the Senate Finance Committee*. Advisory Commission to Study the Consumer Price Index.
- Bureau of Labor Statistics (2012). *Relative Importance of Items in the Consumer Price Index: Tables 1 - 7, Relative Importance of Components in the Consumer Price Index, all areas*. Bureau of Labor Statistics. Retrieved from <http://www.bls.gov/cpi/cpiri2011.pdf>
- Caballero, R. J., & Engel, E. M. (2007). Price stickiness in Ss models: New interpretations of old results. *Journal of monetary economics*, 54, 100-121.

- Calhoun, C. A. (1996). OFHEO house price indexes: HPI technical description. *Office of Federal Housing Enterprise Oversight*.
- Case, B., Pollakowski, H. O., & Wachter, S. M. (1991). On choosing among house price index methodologies. *Real estate economics*, 19(3), 286-307.
- Crawford, M., Mauro, A., & Church, J. (Eds.) (2011). *CPI Detailed Report: January, 2011*. Bureau of Labor Statistics. Retrieved from <http://www.bls.gov/cpi/cpid1101.pdf>
- Crone, T. M., Khettry, N., Mester, L. J., & Novak, J. A. (2011). *Core Measures of Inflation as Predictors of Total Inflation*. FRB of Philadelphia Working Paper No. 11-24.
- Diewert, E. (2003). *The treatment of Owner Occupied Housing and other durables in a Consumer Price Index*. Manuscript, University of British Columbia.
- Diewert, E. (2009). The Paris OECD-IMF workshop on real estate price indexes: conclusions and future directions. *Price and Productivity Measurement Volume 1: Housing*.
- Diewert, W. E., Nakamura, A. O., & Nakamura, L. I. (2009). The housing bubble and a new approach to accounting for housing in a CPI. *Journal of Housing Economics*, 18(3), 156-171.
- Diewert, W. E., & Nakamura, A. O. (2009). Accounting for Housing in a CPI. *Price and Productivity Measurement Volume 1: Housing*.
- Garner, T. I., & Verbrugge, R. (2009). The puzzling divergence of us rents and user costs, 1980-2004: Summary and extensions. *Price and Productivity Measurement Volume 1: Housing*.
- Genesove, D. (1999). *The nominal rigidity of apartment rents* (No. w7137). National Bureau of Economic Research.
- Genesove, D. (2003). The nominal rigidity of apartment rents. *Review of Economics and Statistics*, 85(4), 844-853.
- Giacomini, R., & White, H. (2006). Tests of conditional predictive ability. *Econometrica*, 74(6), 1545-1578.
- Gillingham, R. (1980). Estimating the user cost of owner-occupied housing. *Monthly Labor Review*, 103, 31.
- Gillingham, R. (1983). Measuring the cost of shelter for homeowners: Theoretical and empirical considerations. *The Review of Economics and Statistics*, 65(2), 254-265.
- Gillingham, R., & Lane, W. (1982). Changing the Treatment of Shelter Costs for Homeowners in the CPI. *Monthly Labor Review*, 105, 9.

- Greenlees, J. S. (2006). The BLS Response to the Boskin Commission Report. *textitInternational Productivity Monitor*, 12, 23.
- Greenlees, J. S., & Williams, E. (2009). *Reconsideration of Weighting and Updating Procedures in the US CPI*. US Department of Labor, US Bureau of Labor Statistics, Office of Prices and Living Conditions.
- Steven, W. H., & Stephen, A. B. (1990). Quality Adjustments for Structural Changes in the CPI Housing Sample. *Monthly Labor Review*, 113, 40-42.
- Hoffmann, J., & Kurz-Kim, J. R. (2006). *Consumer price adjustment under the microscope-Germany in a period of low inflation* (No. 652). European Central Bank.
- Katz, A. J. (2009). Estimating dwelling services in the candidate countries: theoretical and practical considerations in developing methodologies based on a user cost of capital measure. chapter, 3, 33-50. *Price and Productivity Measurement Volume 1: Housing*.
- Khater, Same (2012): "What Markets offer the Best Return for Single-Family Rental Investors?" *The MarketPulse*, Core Logic. Vol. 1, Issue 4.
- Lane, W. F., Randolph, W. C., & Berenson, S. A. (1988). Adjusting the cpi shelter index to compensate for effect of depreciation. *Monthly Labor Review*, 111, 34.
- Lang, C. (2012). The Dynamics of House Price Capitalization and Locational Sorting: Evidence from Air Quality Changes. *US Census Bureau Center for Economic Studies Paper* No. CES-WP-12-22.
- Nagaraja, C. H., Brown, L. D., & Wachter, S. M. (2010). *House price index methodology*. Working paper.
- Poole, R., Ptacek, F., & Verbrugge, R. (2005). Treatment of owner-occupied housing in the cpi. *Federal Economic Statistics Advisory Committee (FESAC) on December, 9, 2005*.
- Price Statistics Review Committee. (1961). The Price Statistics of the Federal Government. *Review, Appraisal, and Recommendations (Washington, DC, 1961)*, 79.
- Savage, H. A. (1998). *What We Have Learned About Properties, Owners, and Tenants from the 1995 Property Owners and Managers Survey* (Vol. 3). Census Bureau, US Department of Commerce, Economics and Statistics Administration.
- Schultze, C., & Mackie, C. (Eds.). (2001). *At What Price?: Conceptualizing and Measuring Cost-of-Living and Price Indexes*. National Academies Press.

- Shiller, R. J. (1991). Arithmetic repeat sales price estimators. *Journal of Housing Economics*, 1(1), 110-126.
- Sirmans, G. S., MacDonald, L., Macpherson, D. A., & Zietz, E. N. (2006). The value of housing characteristics: a meta analysis. *The Journal of Real Estate Finance and Economics*, 33(3), 215-240.
- Shimizu, C., Nishimura, K. G., & Watanabe, T. (2008). *Residential Rents and Price Rigidity: Micro Structure and Macro Consequences*. Working Paper.
- Shimizu, C., Diewert, W. E., Nishimura, K. G., & Watanabe, T. (2012, May). The Estimation of Owner Occupied Housing Indexes using the RPPI: The Case of Tokyo. In *Meeting of the Group of Experts on Consumer Price Indices Geneva* (Vol. 30).
- Smith, D. A. (1975). *The Flow of Services Approach to Estimating the Homeownership Component of the CPI*. Bureau of Labor Statistics manuscript.
- Stock, J. H., & Watson, M. W. (2007). Why has US inflation become harder to forecast?. *Journal of Money, Credit and Banking*, 39(s1), 3-33.
- Taylor, John (2007): "Housing and Monetary Policy". Presented at the Policy Panel at the Symposium on Housing, Housing Finance, and Monetary Policy
- Taylor, J. B. (2009). *The financial crisis and the policy responses: An empirical analysis of what went wrong* (No. w14631). National Bureau of Economic Research.
- Verbrugge, R. (2008). The puzzling divergence of rents and user costs, 1980-2004. *Review of Income and Wealth*, 54(4), 671-699.
- Verbrugge, R. (2004). The puzzling divergence of aggregate rents and user costs, 1978-2001. In *SSHRC Conference on Price Index Concepts and Measurement, Vancouver* (Vol. 30).
- Verbrugge, R. J. (2011). Do the Consumer Price Index's Utilities Adjustments for Owners Equivalent Rent Distort Inflation Measurement?. *Journal of Business & Economic Statistics*, 30(1), 143-148.
- Verbrugge, R., & Garner, T. I. (2009). Reconciling user costs and rental equivalence: Evidence from the US consumer expenditure survey. *Journal of Housing Economics*, 18(3), 172-192.