Gimme Shelter!
Rents Have Risen, Not Fallen, Since World War II

BY LEONARD NAKAMURA

Two recent studies have concluded that for roughly four decades the measure of inflation for rents in the U.S. consumer price index was substantially underestimated. Why should this mismeasurement be of concern? In this article, Len Nakamura explains that rents are important in measuring the price of housing services for homeowners as well as renters. They are also the main standard against which market participants and others weigh the reasonableness of house prices. In addition, such mismeasurement affected the estimated rate of overall inflation faced by U.S. households during this historical episode.

Measuring rental inflation accurately is important because rents are the largest component in the U.S. consumer price index, representing fully one-third of the consumption basket. This might seem surprising, since the U.S. is largely a nation of homeowners, not renters. However, government statisticians use rents as a proxy for the price of housing services consumed by owners, for reasons we explore in this article. A related reason for accurately measuring rental inflation is that rents represent the main standard against which to measure the reasonableness of house prices.

Len Nakamura is an economic advisor in the Research Department of the Philadelphia Fed. This article is available free of charge at www.philadelphiafed.org/econ/br/index.

According to official U.S. data, while all prices have been rising, rents have been rising more slowly than other prices. From 1942 to 2003, rents, as measured by the U.S. consumer price index, went up less than ninefold, while the consumer price index, excluding shelter, went up more than 10-fold. Thus, the ratio of rents to other prices is 20 percent lower than it was in the 1940s. However, this relative decline took place roughly between 1942 and 1985 — a period during which, as two new studies suggest, rental inflation was underestimated. Figure 1 depicts this relative decline in prices by showing the ratio of rents to other prices, excluding shelter, in the falling beaded line.

These studies have concluded that for roughly four decades, from 1942 to 1985, the measure of inflation for rents in the U.S. consumer price index was substantially underestimated. My study with Theodore Crone and Richard Voith finds an annual understatement of 1.4 percentage points for the rental inflation rate, while one by Robert Gordon and Todd vanGoethem argues for an understatement of 1.2 percentage points. Over time, these errors cumulate into large numbers and result in very different long-term relationships between rents and overall inflation in the U.S. Either set of estimates indicates that rents have generally risen faster than other prices throughout the postwar period; our estimates show rents rising relatively about 50 percent rather than falling 20 percent, depicted in the barbed green line.

Because of the large weight of rents in consumption and the substantial size of the bias, the estimated rate of overall inflation faced by U.S. households is visibly affected. One broad measure of the rate of inflation faced by households in the U.S. is the personal consumption expenditure (PCE) deflator. Many economists consider the PCE deflator to be the best overall measure of inflation.[1]

[1] The personal consumption expenditure (PCE) deflator is produced by the Bureau of Economic Analysis as part of its quarterly estimates of gross domestic product from data collected by other agencies. One of those data-collection agencies is the Bureau of Labor Statistics, which is charged with collecting U.S. price data. The PCE deflator is considered better than the consumer price index for two main reasons. First, it is a broader measure of inflation that, in particular, includes more services, such as medical and financial services; second, it is revised historically to be more consistent and to eliminate past errors.

FIGURE 1

Two Measures of Relative Rents*

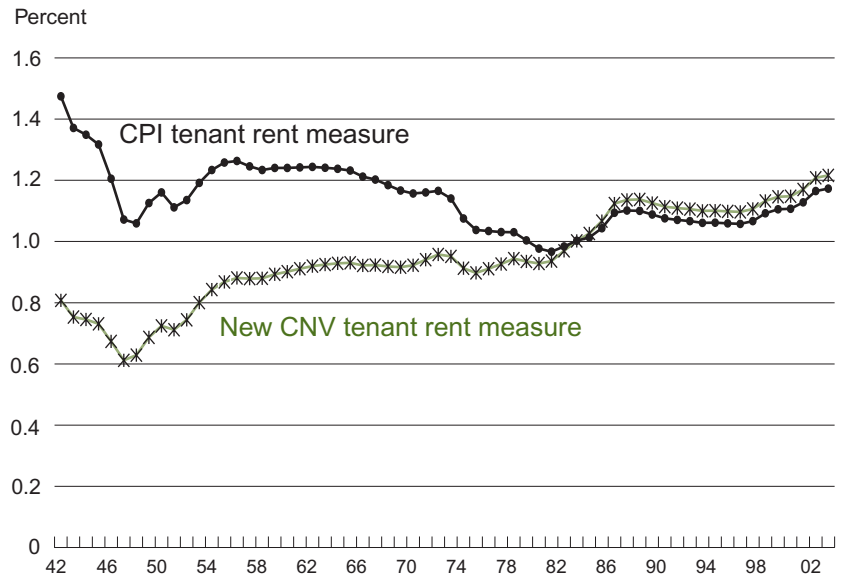

Source: U.S. Bureau of Labor Statistics and Crone et al. (2006).

* Crone, Nakamura, and Voith and Official CPI Measures of Rent Ratio to CPI All Items (excluding shelter)

From 1942 to 1985, the PCE deflator, as currently measured, grew at an annual rate of 4.3 percent. If we use our study's (Crone, Nakamura, and Voith) estimate for rents, that deflator's rate of increase rises to 4.5 percent per year. Similarly, the real growth rate of personal consumption expenditures, as currently measured, grew at 3.8 percent; as revised, it would fall to 3.6 percent. (Real consumption growth per capita would fall from 2.5 percent to 2.3 percent.)

MEASURING INFLATION AND HOUSING SERVICES

How should we measure the part of consumer inflation represented by housing services? In the U.S. consumer price index, produced by the Bureau of Labor Statistics (BLS), and in the personal consumption expenditure deflator, produced by the Bureau of Economic Analysis (BEA), tenant rents are used to measure the price of housing services. These agencies use this method, even though the bulk of housing is occupied by homeowners, who do not pay an explicit rent. To understand why this practice is standard, we need to analyze what the resident of a house consumes. In consumption terms, economists think of housing as providing a service: sheltering the residents and their possessions. This housing service is distinct from the value of the house as an investment. But for homeowners, the house is both a source of housing services and an investment. Therefore, to construct a consumption inflation measure, we have to somehow estimate the value of the housing services consumption component.

Conceptually, a renter and a homeowner get the same housing services, regardless of the form of ownership, if the house is otherwise the same. The renter pays for the service directly. So if we can figure out what the house would rent for, we would know how much the shelter services should cost.[2] If we are lucky, we can find a rental unit just like our house and find out what renters are actually paying landlords. We can then use this as an estimate of the unit's shelter services. Of course, it is often not possible to find rental units precisely equivalent to owner-occupied ones. But since we are interested in the rate of inflation, not the level of prices, rentals that are reasonably similar to the owner-occupied units will be good enough. It is this latter principle that statisticians at the BLS invoke when they measure the housing services of owner-occupied units.

SOME DIFFICULTIES IN MEASURING RENT INFLATION

It would seem that measuring rental inflation should be straightforward, but as so often happens in economic measurement, the details turn out to involve some devilish problems. To measure tenant rents in the U.S., the BLS samples rental properties in urban areas (that is, cities and their surrounding suburbs). Generally speaking, the BLS price inspector obtains this information from the landlord

[2] Housing provides returns to the homeowner in two forms: housing services (or implicit rent) received during the period the homeowner occupies the house, and the value of the house when the homeowner sells it. In turn, the sale value of the house will be derived from the housing services the house provides thereafter.

or real estate manager. If rental units are vacant, their prices are estimated based on the inflation rate at similar units.

Measured Rental Inflation Lags Reality by Three Months. One unusual feature of rents compared with other prices is that the typical rental unit experiences a price change once a year. Thus, if the price inspectors were to check on a given unit every month, 11 times out of 12, the answer would be the rent hasn't changed. So BLS price inspectors collect data on rents only every six months from a given unit. The current monthly rate of rental inflation is then calculated as the average rate of inflation of the units surveyed in that month. One measurement problem shows up immediately: The actual rental price increase at these units could have occurred any time over the past six months, but the increase is included in the index as if it had occurred in the past month.[3] On average, the rental price increase actually measured occurred three months ago; this tends to create a three-month lag in the average time it takes for an increase in rents to show up in the price index. This lag results in rental inflation being understated when it is accelerating and overstated when it slows. If rental inflation is changing rapidly, three months can be a long time, and the measurement error can be significant.

Comparability of Tenant and Owner-Occupied Housing Services. The U.S. consumer price index uses two main measures of housing services: rent of primary residence, for renters, and *owners' equivalent rent* of primary residence, for owner-occupiers. Both are measured using tenant rents. However, owners' equivalent rent inflation differs from rent of primary residence inflation mainly for two reasons. First, owners tend to live in different places and in different types of units than renters. To remedy this imperfect comparability, the BLS gives greater weight to rental units that resemble owner-occupied units, such as single-family detached units, and that are in areas where housing is predominantly owner occupied, such as the suburbs.

Second, for many rental units, landlords directly pay some energy and other utility costs; therefore, these costs are indirectly included in rents paid by tenants. For example, for units that include fuel costs in the rent, rents tend to rise more rapidly when energy costs are rising. Owners, on the other hand, all pay their utility costs directly. To the extent utility costs are included in rents, the BLS has to subtract utility costs from rents to obtain the "pure rents" it needs to calculate inflation for owner-occupied housing.

The BLS has not always used rents to measure the housing services of owner-occupied housing in the consumer price index. From the early 1950s to the early 1980s, it used the so-called acquisition method, measuring house-price inflation, mortgage interest rates, and other out-of-pocket costs of homeownership, such as home insurance. But this method mixed investment returns with housing services' consumption costs, and economists widely viewed it as unsatisfactory. The BEA, which generally uses BLS data to measure inflation and to adjust economic growth for inflation, has consistently measured owner-occupied housing services with rents in the national income accounts as part of its measures of gross domestic product and personal consumption expenditures. Alternative measures of housing services are discussed in *Alternatives to Rent as Measures of Housing Services* and in the working paper by BLS economists Robert Poole, Frank Ptacek, and Randal Verbrugge.

Aging Bias. Aging bias is an additional issue that must be addressed in the data. Does a rental unit remain the same from year to year as it is rented, or does it deteriorate as it ages? If landlords' maintenance and repair activities are not sufficient to keep the average unit as good as new, how important is aging quantitatively? One way to answer this question would be to find two units that are exactly the same but built at different dates. If the older one fetches a lower rent than the newer one, the difference would be attributable to deterioration due to age. But such situations rarely occur and are unlikely to be representative of all units.

Another way to accomplish the same thing is to obtain rental data on a variety of different rental units, along with all of the units' relevant characteristics, and tease out from these data the average impact of aging on rents. The empirical method economists use to do this is called *hedonic regression*. The idea behind this approach is that any product is purchased because of

The U.S. consumer price index uses two main measures of housing services: rent of primary residence, for renters, and *owners' equivalent rent* of primary residence, for owner-occupiers.

[3] The monthly inflation rate is taken to be the monthly rate that would compound to the six-month change. Technically, the monthly log change is calculated by taking one-sixth of the six-month log change.

Alternatives to Rent as a Measure of Housing Services

The Purchase Method. One way to measure owner-occupied housing services inflation is to look at the inflation rate of the houses themselves. We measure the rate of inflation for new cars, refrigerators, and furniture by using the prices of these durable goods, rather than trying to estimate the services we receive from them. So why not do the same for housing? Because, as we have argued, housing prices reflect investment, not just consumption. Moreover, estimating the inflation rate for new homes is actually quite difficult because new homes differ in location as well as in details of construction. In addition, the purchase price of houses doesn't include many of the costs of homeownership, such as taxes and insurance.

The Acquisition Method. As used in the U.S. consumer price index from the mid-1950s to the early 1980s, the acquisition method, also called the asset price approach, included the purchase price of houses, mortgage interest rates, taxes, insurance, and maintenance and repair costs. Among many criticisms of this approach is the fact that the effective cost of a given level of mortgage interest depends on the expected rate of inflation. A mortgage interest rate of 1 percent can be effectively more expensive to the consumer than a mortgage interest rate of 5 percent if the rate of inflation is expected to be negative in the first instance and very high in the second.

The User Cost Method. In this approach, cost depends on the interest and operating costs of the house, less the expected house-price appreciation. At any point in time, we take an individual unit, evaluate its price, and multiply that price by the interest rate less the expected appreciation rate, and add on the operating costs (taxes and maintenance). This is what the owner actually pays to use the unit; it is conceptually the same as the rent, provided there is no risk and no transaction costs. The interest rate calculation should take into consideration the tax treatment of mortgage interest, and the appreciation calculation, the tax treatment of capital gains.

Of course, the expected appreciation rate is never observed directly. Taking into account risks in the valuation of the house together with transactions costs is very difficult. As a consequence, a practical measure of user costs has yet to be set forth.

These difficulties explain why the Bureau of Economic Analysis and the Bureau of Labor Statistics regard tenant rents as the best practical measure of owner-occupied housing services. However, tenant-occupied units and owner-occupied units remain disparate, and as Federal Reserve Board economist Joshua Gallin has shown, the dynamics of owner-occupied house prices and rents are quite different. Thus, further research in this area remains an important item on the price-measurement agenda.

its desirable (or "hedonic") characteristics. For example, a car might have such characteristics as horsepower, gasoline mileage, sunroof, trunk size, interior room, power seats, and so forth. Similarly, a house might have characteristics such as square footage, number of bedrooms and bathrooms, total number of rooms, type of neighborhood, size of garage, central air conditioning, and so forth, as well as age. A hedonic regression would attempt to capture how all of an average unit's characteristics, including age, influenced the unit's rent. If one found that rents fell with age, controlling for changes in other characteristics, it might indicate that, on average, rental units are not maintained in the same condition as when they were first built.

However, using a hedonic regression to estimate the effect of physical deterioration on rents presents two potential problems. The first is the so-called *vintage effect*, which arises when new units have unmeasured quality characteristics that old units do not have. For example, the more extensive use of insulation in houses built after the 1970s may mean that newer houses have higher unmeasured quality — and, thus, fetch higher rents — than older units, but this is not due to the deterioration of older units. On the other hand, if higher quality units remain in the stock of occupied housing while lower quality units are demolished, this may raise the unmeasured quality of older units relative to new units and produce relatively higher rents. But this is not because individual units are getting better over time, just that worse units are disappearing. These so-called vintage effects are hard to separate from the aging effect per se — physical deterioration — on rent.

The second problem in estimating aging's effect on rent is that units of different types (e.g., apartments versus detached houses) may deteriorate at different rates, possibly because the incentives to maintain a unit may differ or maintenance costs may be lower.

In his 1988 articles BLS economist William Randolph took steps to solve

both of these problems in estimating the effect of systematic physical deterioration on rents. Randolph argued that including a sufficient number of housing and neighborhood characteristics in a hedonic equation would render the remaining vintage effect minimal.[4] He included housing characteristics such as the presence of a dishwasher or washer/dryer and neighborhood characteristics such as the percent of the population with a college education. He also estimated different aging effects depending on the number of rooms in the unit, whether the unit was detached, and whether it was rent controlled. His resulting estimate of the *average* effect of aging on rent was -0.36 percentage point a year, meaning that the quality of the average unit deteriorated at that rate. This implied, for example, that if the rental price of an average unit rose 3 percent in a given year, the true rate of rental inflation was 3.36 percent. Since 1988, the BLS has used Randolph's estimating technique, updated over time, to calculate aging's effect, then uses that calculation to adjust the rent component of the CPI by adding on the aging bias. Generally speaking, BLS estimates of the average aging effect have changed very little. In the revised measure of rental inflation developed in our study, aging bias before 1988 is estimated by adopting Randolph's correction of -0.36 percentage point, and raising annual rental inflation rates by 0.36 percentage point.

NONRESPONSE PROBLEMS AND THE BLS CORRECTIONS

Now we turn to the biggest source of error in the historical CPI measures of rent: nonresponse bias. Nonresponse bias is a subcategory of a more general kind of statistical bias: sample-selection bias. Generally speaking, whenever statisticians collect data, they are aware their work is potentially affected by *sample-selection* bias; that is, the data gathered do not reflect the population sampled. This is a problem even for censuses that attempt to reach the entire population, such as U.S. population censuses, which have been shown to have difficulty counting low-income neighborhoods.

For the BLS inflation measures, an important difficulty is ensuring that the price movement of the items surveyed represents the price movement of the items households are actually consuming. For example, the BLS has been criticized for being too slow in introducing new items into its lists of products being priced, such as computers, whose prices decline rapidly when they are first introduced. In that case, the sample-selection bias has tended to cause inflation to be overstated, since prices of older computers fell more slowly than the prices of new ones. The rental sample-selection bias goes in the opposite direction, biasing inflation measures downward: Units where tenants have moved are undercounted, and rents rise faster at these units. Let's look at this problem in more detail.

Before 1942, the BLS gathered rental data primarily from landlords and real estate managers, as it does now. However, during World War II, rent controls were imposed across the nation. As a result, there was concern that landlords and real estate managers might not report rents accurately if they had illegally increased rents. In fact, later analysis suggested that in many neighborhoods, there were more rent increases than had been authorized. So the BLS instructed rental price inspectors to ask tenants for the data and that is what they did from 1942 until the end of 1977.

This meant going to tenants and getting them to agree to an interview and inspection of their dwellings and to answer follow-up mail questionnaires every three months. (In the 1950s the frequency of the mail questionnaires was reduced to every six months.) When tenants moved, the BLS would have to find the unit's new tenants and get them to agree to participate, and the BLS would also have to reinspect the unit to see if the landlord or manager had made any changes to it. If the new tenant could not be contacted soon enough, or if the unit remained vacant in the price collection month, the BLS would not record information on rent at that unit, and any price increase at the unit would be lost. But tenants, when they move, usually move at the time of the unit's annual rent increase. The data lost when tenants move have a much higher probability of including a price increase than the data for a typical unit; therefore, this problem of nonresponse biases the rate of inflation downward. This *nonresponse bias* problem was revealed in a study by two BLS economists, Joseph Rivers and John Sommers, and my study with Crone and Voith.

When a rental unit changes tenants, there is a tendency for the price increase to be greater than if the tenant had stayed.

[4] Gordon and vanGoethem argue that Randolph's methodology insufficiently accounts for quality improvements in housing. For example, Randolph's methodology will not capture the change in quality if homes are constructed with more thermal insulation.

Compounding this issue is another interesting problem revealed by the BLS data: When a rental unit changes tenants, there is a tendency for the price increase to be greater than if the tenant had stayed. Rents for units whose tenants change rise about one-third faster than rents for tenants who continue in residence. So not only were some rent increases lost, but the ones lost were generally larger. One possible reason for this phenomenon is that when landlords raise the rent too high, tenants leave. But that doesn't explain why the rental inflation rate tends to be low if the tenant continues to stay in the apartment, at least for a few years, or so the paper by Hebrew University professor David Genesove argues. Instead, he suggests, finding a good tenant is not always easy for landlords, and so landlords tend to keep the rent low for continuing tenants.

Many of these problems were solved in 1978, when the BLS made a major revision to the methods by which it collects the data for the consumer price index. The 1978 revision, perhaps the BLS's most expensive makeover ever of its consumer price index, was intended to place the consumer price index on as sound a statistical footing as possible. As part of this revision, the BLS elected to shrink the size of the rental sample but put more resources into obtaining high response rates from the units. One step was to permit the price inspectors to go back to surveying landlords and real estate managers as well as tenants; in practice, this meant surveying mainly landlords and managers.

The BLS also conducted a number of studies examining the impact of the revised methods. The paper by Rivers and Sommers was one result; they found that new tenants were indeed now being included in the survey and that the rents for these new tenants tended to reflect higher rates of inflation than other units. They also pointed out that the new method still omitted price increases in units that remained vacant at the time of their price inspection, and they were able to show that this produced a continuing downward bias in the price index.

Their work also pointed up a second problem: *recall bias*. One part of the 1978 revision to rent increases was a very clever idea: asking whether the rent had increased in the past month. As pointed out earlier, using the six-month change as if it had occurred in the past month creates a lag in the data. In the 1978 revision, the BLS began using a weighted average of the past month's increase with the six-month increase to create a more current index. Indeed, the way the BLS planned to do this would almost completely eliminate the lag in the index. Unfortunately, it turned out that recall of rent increases in the past month was very poor, perhaps because respondents perceive the rent to increase when the former tenant moves out, rather than at the start of the new tenant's occupancy (the BLS's definition).[5] Whatever the reason, adding the one-month rent increases created an additional downward bias. So even after 1978, there continued to be downward biases, which were only fixed at the beginning of 1985.

NEW MEASURES OF RENTS

From 1942 to 1985, primarily because of nonresponse bias, the rent measure was understated. But we have direct evidence of the size of the bias only for 1978 forward. What to do? Our study attempts to "backcast" the size of the bias by setting up a model of the BLS measurement process, including various measured characteristics of the housing market, such as how often tenants move, how long apartments are vacant, and how much rents rise when tenants move. Figure 2 summarizes the contributions of nonresponse bias, as measured by our model, aging bias, and recall bias to our new estimate of rental inflation from 1942 to 1985.

Most of the aspects of this model are testable using BLS data on rents. Such a data set, from the period 1985 to 1988, was made available by Genesove, who had used it in his study of rent dynamics. My coauthors and I were able to show that our model would have given a good approximation of biases from 1985 to 1988, a period of relatively low inflation, even though most of the parameters were calculated based on data from the high-inflation period of the late 1970s and early 1980s. We predict, for example, how much the inflation rate should change if tenants who move are omitted from the sample and then we check whether the rate changes by that proportion in Genosove's data. That parameters taken from a high-inflation episode can be used to match data from a low-inflation episode provides some assurance that the model can be used across periods that include episodes of both types.

We then used the model to estimate that the CPI tenant rent inflation from 1942 to 1985 was too low by 1 percentage point annually because of nonresponse bias and recall bias. During this time, it appears that the BLS missed nearly one out of three rental increases. Given that an aging bias of nearly 0.4 percent was also present in these data, we conclude that the CPI for tenant rents was downwardly biased by 1.4 percentage points annually for

[5] Recall bias was somewhat worse when landlords and managers were the respondents; since landlords and managers have good records upon which to base their answers, this suggests a conceptual confusion on the part of the respondents rather than a factual one.

more than four decades.

This is a very large bias, cumulated over time. As noted earlier, we find that rather than falling 20 percent, as in the BLS estimates, rents rose 50 percent relative to other prices from 1942 to 1985. A similarly large downward bias estimate has also been put forward by Gordon and vanGoethem, using entirely separate data.

A Second Measure of Rental Inflation. Our approach was to take BLS rental information, based on information for individual housing units over time, and correct its biases. An alternative way to measure inflation – the route taken by Gordon and vanGothem – is to measure the average cost of all rental units at two different dates and ask how much the quality of the average unit changed between the two dates. For their purposes, the U.S. censuses of housing and population, conducted every decade, and the American Housing Surveys, conducted at first annually and now every two years, provide estimates of average rents going back to 1930. These censuses and surveys also provide data on various features of the housing units, such as the presence of indoor plumbing, central air conditioning, and so forth.

Unfortunately, as we go back in time, the censuses provide fewer details. The earliest census Gordon and vanGoethem use has very little in the way of detail. Somewhat more detail on rental characteristics is available from a study by Clair Brown that uses budget studies performed by the Bureau of Labor Statistics going as far back as 1918. To estimate quality with skimpier data, Gordon and vanGoethem analyzed more recent decades to put reasonable bounds on their estimates for earlier periods. The main quality adjustments are based on trends in plumbing, central heating, and electrification. They can test their estimates by examining how far off they would be if they used them on more recent data; by and large, their estimates seem to be reasonable.

This admittedly crude methodology does appear to make the best possible use of data other than the consumer price index data we use. Gordon and vanGoethem indicate that from 1940 to 1985, annual inflation has been understated by 1.2 percentage points; this is not far from our study's estimate of 1.4 percentage points.

Gordon and vanGoethem confirmed their estimate by looking at information on rents from Evanston, Illinois, from 1925 to 1999. They used classified advertisements on rents from a local newspaper to construct a rental price index. They were able to collect not only rents, along with some information on quality (such as number of rooms), but also matching apartment rents at specific addresses, which is close to the BLS procedure. They constructed two rental indexes using these two types of data and found that they broadly agree. From 1940 to 1985, rents in Evanston, Illinois, rose roughly 1.6 percent faster than CPI rent inflation. However, one might worry that Evanston, a relatively wealthy suburb, has done better than the average location in the United States.

Taken together, these studies paint a very broad picture of inflation bias in rents from 1940 to 1985. Two very different approaches find bias in the same direction and of the same general magnitude for this period.

CONCLUSION

Housing services are an important part of what we consume. As we have seen, measuring inflation in housing services has, in the past, raised chal-

FIGURE 2

Components of Estimated Rental Inflation, 1942 to 1985

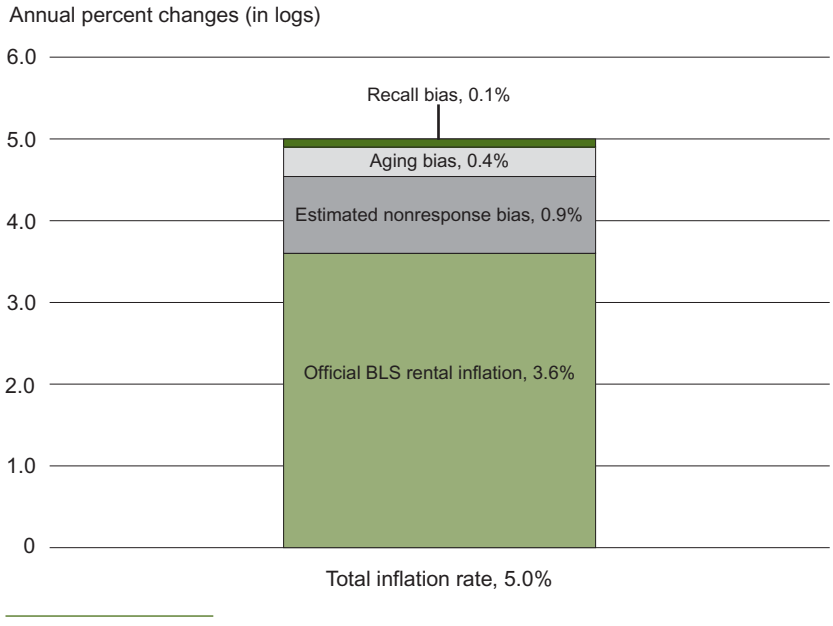

Source: U.S. Bureau of Labor Statistics and Crone et al. (2006)

lenging problems that have not always been immediately recognized. One consequence is that our historical record of rents appears to be inaccurate. At the same time, we must recognize that this is an area in which the Bureau of Labor Statistics took vigorous steps to improve its measures. As a consequence, many of the problems that affected this measure have been solved.

We have argued that measuring rents accurately is important because housing services are a large part of consumption. Another reason accurately measuring rents matters is that since housing services are the benefit we receive from homeownership, rents are an important measure of house values. Inaccurate data on rents may generate conundrums when economists and others seek to understand house prices.

In a recent article in which he argues that house prices are now too high, Yale economist Robert Shiller points out that since 1913, housing prices have risen relative to other prices, while rents have fallen. This is a puzzle because, over long periods, one might expect prices and rents to move together, since rents provide the economic basis for house prices. In fact, once we adjust for nonresponse and aging bias, both rents and house prices have risen over the past 90 years; this provides one possible solution to Shiller's puzzle.

In addition, having more accurate historical inflation statistics helps us better understand our economy. Many economic propositions depend on statistical models that can be accurately measured only with long data series. By improving this important economic series, we improve the ability of the economics profession to sort out good theories from bad ones.

Finally, economists have puzzled over the productivity slowdown that the U.S. experienced from 1975 to 1995, when output per hour in the nonfarm business sector rose just 1.5 percent annually, compared with 2.4 percent annually from 1955 to 1975. If inflation has been understated, as I have argued it has for rents, then real output growth will tend to be overstated because for a given level of nominal rent payments, lower prices imply higher real consumption. While our new data do not make the post-1975 productivity slowdown vanish, they do reduce its size. In particular, our new data argue that output per hour was actually a bit lower from 1955 to 1975, growing 2.2 percent annually, the postwar average. However, the growth rate from 1975 to 1995 is also somewhat lower, at just 1.4 percent a year, since the data bias continued until 1985.

A difficult question that we have not fully faced up to in this article is: How accurately do rents for tenant units – even when adjusted as the BLS does – capture the housing services of owner-occupied units? This is an important area for future research.

REFERENCES

Brown, Clair. *American Standards of Living, 1918-1988*. Cambridge: Basil Blackwell, 1994.

Chatterjee, Satyajit. "Homeownership, Taxes, and the Allocation of Residential Real Estate Risks," Federal Reserve Bank of Philadelphia *Business Review* (September/October 1996).

Crone, Theodore M., Leonard I. Nakamura, and Richard Voith. "The CPI for Rents: A Case of Understated Inflation," Federal Reserve Bank of Philadelphia Working Paper 06-7 (January 2006).

Crone, Theodore M., Leonard I. Nakamura, and Richard Voith. "Hedonic Estimates of the Cost of Housing Services: Rental and Owner-Occupied Units," in W. Erwin Diewert, Bert M. Balk, Dennis Fixler, Kevin J. Fox, and Alice O. Nakamura, eds., *Prices and Productivity Measurement.* Trafford Press, forthcoming.

Davis, Morris E., and Jonathan Heathcote. "The Price and Quantity of Residential Land in the United States," Working Paper (October 2005).

Gallin, Josh. "The Long-Run Relationship Between House Prices and Rents," Federal Reserve Board FEDS Working Paper 2004-50 (September 2004).

Genesove, David. "The Nominal Rigidity of Apartment Rents," *Review of Economics and Statistics* (November 2003).

Gordon, Robert, and Todd vanGoethem. "Downward Bias in the Most Important CPI Component: The Case of Rental Shelter, 1914-2003," in Ernst R. Berndt and Charles M. Hulten (eds.), *Hard to Measure Goods and Services: Essays in Honor of Zvi Griliches*. Chicago: University of Chicago, forthcoming.

Haines, Michael R., and Allen C. Goodman. "A Home of One's Own: Aging and Homeownership in Late 19th Century and Early 20th Century America," in David I. Kertzer and Peter Loslett, eds., *Aging in the Past: Demography, Society, and Old Age*. Berkeley, CA: University of California Press, 1993.

Nakamura, Leonard I., "Measuring Inflation in a High-Tech Age," Federal Reserve Bank of Philadelphia *Business Review* (November-December 1995).

Nakamura, Leonard I. "The Retail Revolution and Food-Price Mismeasurement," Federal Reserve Bank of Philadelphia *Business Review* (May-June 1998).

Poole, Robert, Frank Ptacek, and Randal Verbrugge. "Treatment of Owner-Occupied Housing in the CPI," BLS Working Paper at: www.bls.gov/bls/fesacp1120905.pdf.

Randolph, William C. "Estimation of Housing Depreciation: Short-Term Quality Change and Long-Term Vintage Effects," *Journal of Urban Economics*, 23 (1988a), pp. 162-78.

Randolph, William C. "Housing Depreciation and Aging Bias in the Consumer Price Index," *Journal of Business and Economic Statistics*, 6 (1988b), pp. 359-71.

Rivers, Joseph D., and John P. Sommers, "Vacancy Imputation Methodology for Rents in the CPI," *Proceedings of the ASA Economics and Business Section* (1983), pp. 201-05.

Shiller, Robert. "Long-term Perspectives on the Current Boom in Home Prices," *Economists' Voice* (March 2006).

Sinai, Todd, and Nicholas S. Souleles. "Owner-Occupied Housing as a Hedge Against Risk," Federal Reserve Bank of Philadelphia Working Paper 05-10 (2005).