

A Panel of Interarea Price Indices for All Areas in the United States 1982-2012

Online Appendices

Paul E. Carrillo
Department of Economics
George Washington University
2115 G Street, NW
Washington, DC 20052
pcarrill@gwu.edu

Dirk W. Early
Department of Economics and Business
Southwestern University
Georgetown, TX 78627
earlyd@southwestern.edu

Edgar O. Olsen*
Department of Economics
University of Virginia
Charlottesville, VA 22901
Phone: 434-924-3443
Fax: 434-982-2904
eoo@virginia.edu
*Corresponding author

December 20, 2013

Table of Contents

Appendices

A	Existing Interarea Price Indices	1
B	Differences in Relative Interarea Housing Prices for Units of Different Qualities	9
C	Alternative Housing Price Indices Based on CSS Data	10
D	Approximations of Ceiling Rents Used to Estimate Truncated Regression Model	14
E	Construction of Price Indices for Other Goods and All Produced Goods	17
F	Explanatory Variables in Regression Explaining Differences in Nonhousing Prices	24
	Appendix References	26

Appendix A

Existing Interarea Price Indices

This appendix describes existing interarea price indices and the data that underlie them. It discusses the strengths and weaknesses of the alternative data sources and compares their housing data with ours. We consider separately sources that collect data on the prices of a wide range of goods and services and those that can be used to produce only housing price indices.

Except for housing price indices, most interarea price indices used in economic research have been based at least in part on data collected by the Bureau of Labor Statistics (BLS) for the production of its time-series consumer price indices (hereafter CPI data) or the American Chambers of Commerce Researchers Association (ACCRA), now the Council for Community and Economic Research (C2ER). Interarea housing price indices are sometimes based on data from these sources, but more often on data from the American Housing Survey (AHS), American Community Survey (ACS), or Decennial Census. Between 1967 and 1981, the BLS used the CPI data to produce cross-sectional price indices for 6 broad categories of goods and an overall consumer price index for 39 metropolitan areas and the non-metropolitan urban areas in four regions.¹ Since the BLS discontinued this series, analysts at the BLS and Bureau of Economic Affairs (BEA) have produced exploratory interarea price indices for various years and levels of geography using the CPI data, sometimes supplemented with data from the ACS or Decennial Census. Because access to the CPI data is restricted to government employees and the exploratory price indices have limited geographic and temporal coverage, others whose research would benefit from interarea price indices have relied almost exclusively on other sources, primarily the ACCRA nonhousing price indices and data from the AHS, ACS, or Decennial Census to create housing price indices. ACCRA has produced a series of cross-sectional interarea price indices for many urban areas since 1968. For many years, it has collected price data from more than 300 urban areas each quarter. Neither the CPI nor the ACCRA data are collected from all localities, and some analysts have predicted price indices for areas where CPI and ACCRA data are not collected.

¹ Johnson, Rogers, and Tan (2001, pp. 32-33) provide an account of the development and demise of the BLS price indices.

Overall Consumer Price Indices

Because the ACCRA or CPI data underlie almost all research that accounts for interarea differences in nonhousing as well as housing prices, it is important to describe them in some detail and discuss their strengths and weaknesses. Since our overall price index depends in part on the ACCRA price indices for nonhousing goods and this is the best publicly-available interarea price index produced on a regular schedule for many places, it receives special attention.

For many years, ACCRA and its successor C2ER have published an overall consumer price index and six composite price indices that are expenditure-weighted averages of price indices for 59 categories of goods.² They are a series of cross-sections rather than a panel. ACCRA picks one narrowly defined good, for example, 160-count Kleenex brand facial tissue, to represent price differences for all goods in a category. That is, it assumes that if the particular good priced is 10 percent more expensive in one location than in another, all goods in its group are 10 percent more expensive. Price indices are produced quarterly for urban areas that account for about 70 percent of the U.S. urban population. The Statistical Abstract of the United States has reported these price indices since 1990, and C2ER sells an electronic file with the prices of the individual goods and services underlying the indices.

The primary concerns about the ACCRA price indices have been the small number of price quotes in each area (5 per quarter for each good), volunteer data collectors, and expenditure weights applicable to households in the top quintile of the income distribution with a member in a professional or managerial occupation. Except for housing, the narrow definition of the goods involved ameliorates the objection based on small sample size. There is surely less variation in the prices of narrowly defined goods than more heterogeneous goods. No direct evidence shows that ACCRA's volunteers are less accurate than professionals in recording price data. ACCRA provides its volunteers with detailed instructions, and it reviews their reported prices carefully for seeming anomalies (Council for Community and Economic Research, 2006, pp. 1.4-1.5). Because ACCRA reports the individual prices that underlie its overall consumer price index, alternative expenditure weights can be used to produce an overall price index and price indices

²Council for Community and Economic Research (2006) documents their data collection procedures and price index construction.

for composite commodities such as food. Koo, Phillips, and Sigalla (2000, pp. 130-131) find that replacing ACCRA's expenditure weights with weights reflecting average expenditure shares has very little effect on the overall price index. In section 5, we report similar results for our new price indices.

The ACCRA housing price index leaves much more room for improvement than its other price indices. The main problems are accounting for differences in housing and neighborhood characteristics and predicting the rental value of owner-occupied units. The data set underlying our housing price is much superior to the ACCRA data set in these regards, and our sample size is much larger (170,000 versus 3,000).

Accounting for the many differences in the characteristics of the dwelling unit and its neighborhood is a perennial problem in constructing an accurate interarea housing price index. Differences in these characteristics lead to enormous differences in the rental value of dwelling units within a given market, and the average values of these characteristics are not the same across markets. ACCRA does account for many such differences. Its housing price index is a weighted average of a price index for homeowners and renters. For both, ACCRA controls accurately for the size of the unit and (for homeowners) the size of the parcel. To control for the condition of the unit, ACCRA prices apartments less than 10 years old whenever possible. For homeowners, it prices newly built units. A much greater attempt is made to account for amenities for homeowners than renters. For renters, ACCRA makes no direct attempt to account for amenities beyond the provision of a stove and refrigerator. To account for the many differences between units that are not directly specified, ACCRA attempts to price units occupied by managerial and professional couples in the top fifth of the income distribution. The range of differences in the overall desirability of units among this group is certainly much less than for the entire population. Nevertheless, the remaining differences in the characteristics of the structures and their neighborhoods among the units in the ACCRA sample might be significant. This makes ACCRA's small sample size in each area (5 rental and 5 owner) more problematic for housing than for other goods.

A second problem with the ACCRA housing index is prediction errors in the price index for homeowners. This is particularly important because the housing price index for homeowners accounts for 80 percent of the overall housing price index and 24 percent of the overall consumer price index. Our purpose is to produce a housing price index that compares the cost of

occupying an identical unit during a year in different locations. The rents of apartments correspond exactly to this concept. For homeowners, the ideal is how much their units would rent for. ACCRA's homeownership price index can be viewed as an approximation of this ideal. ACCRA attempts to determine the sales price of very similar new houses in all locations. It then determines the level payment on a 30 year mortgage with a 25 percent downpayment at the average local mortgage interest rate. The average of these level payments across all units in a locality scaled to have a mean of 100 across all localities is the housing price index for homeowners. The question is how well these level payments reflect the market rent of the unit during its first year, or more precisely whether they are proportional to the market rent of these units across areas. The sales price of a house depends not only on its rental value during the current year (net of depreciation and operating expenses) but also its expected net rental value in future years. The ratio of current rent to sales price for identical units can be different in different locations due to different expectations about the future. For example, suppose that it is announced that a large plant will be constructed in a small community in several years. This would have an immediate effect on the sales prices of existing houses and vacant land, but it would not affect current rents. Winters (2009) finds that an index of the sales price of identical houses across areas is a poor proxy for an index of the rents of identical units.

Due to concerns about the accuracy of the ACCRA housing price index (namely, its failure to account sufficiently for differences in the characteristics of dwelling units and their neighborhoods and inaccuracies in predicting the market rent of owner-occupied units) and the availability of alternative data sets for producing such indices, many studies that have used the ACCRA price indices for nonhousing goods have created alternative housing price indices based on the ACS or Decennial Census. Applications that account only for differences in housing prices rarely use the ACCRA housing price index. Instead, they use data from the AHS, ACS or Decennial Census to create their own index [Albouy, 2009; Chen and Rosenthal, 2008; Follain, 1979; Gabriel and Rosenthal, 2004; Malpezzi, Chun, and Green, 1998; and Moretti, 2013].

From time to time, analysts at the BLS and BEA have used the data set underlying the time-series CPI to produce exploratory cross-sectional price indices for broad categories of goods and an overall consumer price index. Unlike ACCRA, the CPI data set is collected by professionals. It also has more individual price observations each year than ACCRA (about

1,000,000 versus 360,000) and prices many more goods (about 370 versus 59).³ Like ACCRA, the CPI data set covers only urban areas. However, the CPI collects data from many fewer urban areas than ACCRA (87 versus more than 300), albeit selected by stratified random sampling to represent all urban areas [Moulton, 1995, pp. 183-184]. The housing information in the CPI comes from a survey of about 50,000 dwelling units. This is much larger than the ACCRA survey (about 3,000) and much smaller than ours (about 170,000). In some years, the BLS housing survey has contained owner-occupied as well as rental units. In other years, it has been limited to rental units. However, in all years since 1982, its housing price index has been based in part on estimates of the market rental value of owner-occupied units [Ptacek and Baskin, 1996]. The CPI Housing Survey contains only a few rudimentary housing characteristics. BLS and BEA analysts who have used it to produce cross-sectional price indices have typically supplemented it with neighborhood characteristics from the Decennial Census. Unlike the ACCRA data set, the CPI data is not available to independent researchers.

The most important BLS and BEA studies are Kokoski, Cardiff, and Moulton (1994) and Aten, Figueroa, and Martin (2011). Based on CPI data from July 1988 through June 1989, Kokoski, Cardiff, and Moulton produced price indices for 11 categories of goods for 44 areas (32 specific urban areas and all other urban areas divided into 12 categories by region and metro status), but did not produce an overall consumer price index. In order to provide estimates of real income, expenditure, and output at low levels of geography based on BEA's estimates of their nominal magnitudes, Aten, Figueroa, and Martin (2011) used the CPI data combined with data from the 2005-2009 ACS to produce an overall consumer price index for 361 metropolitan areas and 51 states.⁴ The limited geographic coverage of the CPI data set (87 urban areas) necessitates assumptions in getting from the data used to the results. For example, except for 31 specific metro areas, all counties in the same region and with the same metro status are assumed to have the same prices for all nonhousing goods.

Most BLS and BEA studies have better data for the specific urban areas where the BLS collects CPI data than ACCRA. For these areas, the BLS and BEA price indices are almost

³ The ACCRA sample size is now smaller. Since 2007, it has collect data for only the first three quarters of the year. The results reported in the fourth quarter are averages of the previous three quarters.

⁴ The single cross-section reported and used in the analysis presumably reflects the average difference in the overall consumer price index across areas from 2005 through 2009.

surely better than the ACCRA index. For other areas where ACCRA collects data, the opposite may be true.

Koo, Phillips, and Sigalla (2000) shed light on the reliability of the ACCRA index compared with an overall price index based on CPI data, albeit in a comparison limited to 23 metropolitan areas. Specifically, they compare ACCRA's cost-of-living index with a cost-of-living index based on Kokoski, Cardiff, and Moulton (KCM)'s price indices. When the same simple formula and expenditure weights are used to produce the cost-of-living indices and the two indices are rescaled to have the same mean, the mean of the absolute percentage deviations between the cost-of-living indices is 5.8 percent.

Finally, some applications such as Albouy (2012) have expanded the geographical coverage of the ACCRA indices in a single year by predicting price indices for places where data were not collected. Others such as Baum-Snow and Pavan (2012) have created a panel of price indices by applying BLS time-series price indices to a single cross-section.

Housing Price Indices

Many studies have created only interarea housing price indices. Some were devoted exclusively to producing them, and others used them to study specific questions. In some applications, price indices for other goods were not relevant. In others, the underlying model included nonhousing goods but the authors assumed that the prices of these goods were the same everywhere.

The most reliable housing price indices have been produced with data from the metropolitan sample of the American Housing Survey (AHS). In the most recent detailed study, Thibodeau (1995) used these data to produce a panel of interarea housing price indices that accounted for many housing and neighborhood characteristics and paid careful attention to model specification. Two major shortcomings of this panel for many purposes are its vintage and geographic coverage. It is only available for about 44 metropolitan areas (about 11 per year in each year between 1984 through 1992 with each area represented in several years). Blackley and Follain (1986), Follain and Ozanne (1979), and Thibodeau (1989) used AHS data and similar methods to produce housing price indices for selected metropolitan areas in earlier years. Ozanne and Malpezzi (1985) and Kiel and Zabel (1997) concluded that the AHS's biggest drawback is its lack of objective information about the neighborhood and location within the metro area. Almost all of the information about neighborhood conditions comes from asking the

respondents, and no AHS reports location for an area with a population less than 100,000.

Normally, the BLS and BEA studies that produce price indices for many broad categories of goods or an overall consumer price index do not carefully document the methods used to create their housing price index.⁵ However, in a methodological paper devoted to comparing housing price indices based on different statistical models, Moulton (1995) describes in some detail the CPI housing data and the general approach used to create the price indices in most BLS and BEA studies. Like KCM (1994), this paper produces housing price indices for 32 specific urban areas and all other urban areas divided into 12 categories. The CPI housing data has the same shortcomings as the ACCRA data, namely, limited information about housing characteristics and prediction errors in estimating the market rents of owner-occupied units. A comparison of Moulton's Table 1 with our online table A-1 clearly shows that the CPI data set contains many fewer housing characteristics than the data set underlying our housing price index. It contains no information about the condition of the dwelling unit or its convenience to jobs, shopping, or recreation facilities. It also contains data for many fewer geographical areas. Construction of housing price indices from the CPI data set has always involved estimating the rental value of owner-occupied units. At some times, this has been the owner's guess. At other times, it has been based on estimating a simple statistical model [Ptacek and Baskin, 1996, p. 34]. In contrast, our data set is limited to rental units and hence does not involve inaccuracies in predicting the market rents of owner-occupied units.

Malpezzi, Chun, and Green (1998) have produced a housing price index for 1990 for 272 MSAs and the nonmetropolitan areas within each state based on the limited set of housing characteristics in the Decennial Census. Their hedonic equation explaining rent has 19 regressors representing only 11 rudimentary characteristics such as the number of rooms and bedrooms, the existence of complete plumbing and kitchen facilities, and the age of the structure. The ACS has essentially the same housing characteristics. Dwelling units that are the same with respect to these characteristics can differ enormously in their condition, amenities, neighborhoods, and convenience to jobs, shopping, and recreation facilities. If there were differences in the mean values of these omitted characteristics across areas among units with the

⁵ This is understandable. Although the hundreds of goods in the CPI survey are very narrowly defined, they are not completely homogeneous. The survey collects data on differences in at least a few characteristics of most goods, and the BLS and BEA analysts estimate hedonic equations for most to produce a price index for that good. So the housing hedonic equation is only one of many involved in their analysis.

same values of the included characteristics, their housing price index would be biased on that account.

In addition to the preceding studies whose primary purpose was to produce interarea housing price indices, some applications such as Albouy (2009), Chen and Rosenthal (2008), and Moretti (2013) have used data from a decennial census to produce such indices for specific years and places to study particular questions. Given the underlying data, these are at least somewhat cruder than those developed in this paper.

Appendix B

Differences in Relative Interarea Housing Prices for Units of Different Qualities

The housing price indices produced by Thibodeau (1989, 1995) shed light on the differences between housing price indices for units of different qualities. Thibodeau produced separate rental housing price indices for units built in the last three years, older standard units, and older substandard units. Less than 4 percent of rental units were built in the last 3 years and less than 7 percent were severely or moderately inadequate according to the AHS's definitions of these terms. Because he used data from the Metropolitan AHS, he had sufficient observations to estimate separate hedonic equations for each place and year. He used these estimated hedonics to predict the market rent of units in each time and place at the national average values of the regressors for the three categories of units. To compare the extent to which these predicted market rents indicate the same percentage differences in housing prices, we first rescale the three price indices in each study to have a mean of 1. After deleting the obviously erroneous result reported for Indianapolis's first survey in the 1989 study, there were 163 observations for each price index in the 1989 study and 103 in the 1995 study. For the observations in the 1989 study, the correlation between the price index for new units (PNew) and the price index for older standard units (PStand) is .94 and the correlation between PStand and the price index for older substandard units (PSub) is .98. The mean of the absolute percentage deviations between PNew and PStand is 7.6 percent and between PSub and PStand is 4.8 percent. For the observations in the 1995 study, the correlation between PNew and PStand is .90 and the correlation between PStand and PSub is .96. The mean of the absolute percentage deviations between PNew and PStand is 9.3 percent and between PSub and PStand is 6.5 percent.

Appendix C

Alternative Housing Price Indices Based on CSS Data

To check the robustness of the results, alternative methods were employed to produce housing price indices based on the CSS data. This section describes these methods, and it compares the alternative price indices with the basic index. Online table A-2 reports the results of the hedonic regressions and online table A-3 the corresponding price indices. For each alternative housing price index, online table A-5 reports the results of OLS estimation of a linear regression of the alternative price index on the basic index, after scaling each so that its mean is 1. It also reports the mean and maximum absolute percentage difference between alternative price indices across all areas.

If the price indices were identical, the slope coefficient and coefficient of determination would be 1. The null hypothesis for testing the proportionality of the price indices on average is that the slope coefficient is 1. Because the price indices are scaled so that their means are one, the estimated constant term is one minus the estimated slope, and the test of the hypothesis that the slope is equal to one yields the same conclusion as the test of the hypothesis that the intercept is zero. For this reason, we report only the estimated slope coefficient and its standard error. Although we can reject the proportionality hypothesis at the usual levels of significance in most cases, the magnitudes of the deviations from proportionality are minuscule in all cases.

Median Regression Estimation

Our first alternative housing price index is based on estimating the coefficients in the regression model by minimizing the sum of absolute deviations, the usual estimator of the median regression model. This tests the sensitivity of our price index to assumptions about the conditional distribution of the error term in the hedonic regression and provides a reasonable alternative estimator of its parameters under the standard assumptions. Online table A-5 shows that the slope and coefficient of determination deviate only slightly from one, the mean absolute percentage difference between the price indices across all areas is only one percent, and the maximum absolute percentage difference is less than six percent.

Alternative Treatments of Missing Values

Including missing value indicators allowed nearly all observations to be used in estimating the hedonic regression and constructing the basic housing price index. An alternative method is to omit observations with missing values for any variables, normally called a complete case analysis (CCA). This requires the omission of roughly half of all observations. In addition to a CCA based on the full set of variables, a CCA based on a shorter list of variables, omitting those variables with the most missing values, was also employed. The second and third rows of online table A-5 report the comparisons of the price indices based on these regressions with our basic housing price index. In both comparisons, the slope and coefficient of determination deviate only slightly from one, the mean absolute percentage difference between the price indices is less than one percent, and the maximum absolute percentage difference is less than six percent.

Truncated Regression

Two of the three HUD programs of tenant-based housing assistance in 2000, the year of our data, had ceiling rents. On October 1, 1999, HUD began to phase out its old Section 8 certificate and voucher programs in favor of the new Section 8 Housing Choice Voucher Program.⁶ This transition continued into 2002. About 90 percent of the households in our sample received assistance under the old housing certificate or new housing voucher program. Unlike the old voucher program, these programs have upper limits on the rent of the unit that can be occupied. To the extent that these limits are binding constraints for some voucher recipients, they imply that voucher recipients, especially those who occupy housing that is the best with respect to observed characteristics, will occupy units that are worse on average than other units with the same observed characteristics. This leads to bias in OLS estimators of the hedonic equation. The extensive set of housing and neighborhood characteristics included in the hedonic regression reduces the variance in its error term and hence ameliorates this problem. However, in addition to price indices based on the standard OLS estimation of the hedonic equation, we produce a price index based on maximum likelihood estimation of a stochastic model that accounts for this truncation, and we compare this price index with the price index based on the standard OLS

⁶ Olsen (2003, pp. 400-404) describes the main features of these programs and how they affect the budget spaces of families offered these subsidies.

estimation. Estimation of this model requires information on the ceiling rent that faced each certificate and voucher recipient in our sample. The CSS data does not include these ceiling rents. Using the approximations described in Appendix D, we estimated a hedonic regression model based on standard truncated regression assumptions [Maddala, 1983]. The fourth row of online table A-5 shows that the resulting housing price index differed little from the basic index. The slope and coefficient of determination deviate only slightly from one, the mean absolute percentage difference between the price indices across all areas is less than two percent, and the largest absolute percentage difference is fifteen percent.

Separate Hedonics for Each Geographic Area

Finally, Early (2006) has produced housing price indices with the CSS data based on the estimation of separate hedonic equations in each location, albeit with a parsimonious set of explanatory variables to accommodate the small sample sizes in some areas. Precise estimation of the mean or median rent of units with specified characteristics in an area requires a substantial sample size relative to the number of characteristics involved in the hedonic regression.⁷ The CSS data set has a relatively small number of observations in some metropolitan areas and the nonmetropolitan parts of some states. To retain the maximum number of observations, Early imputed the missing values of explanatory variables using Stata's imputation procedure. To increase the sample size in each area, he combined CSS data for three years and included year dummy variables in the hedonic regression. Even with these methods for expanding the sample size, 21 areas did not meet his low cutoff of 110 observations for estimation of the hedonic equation. Eighteen small metropolitan areas were combined with the nonmetropolitan part of their states and price indices were not produced for the nonmetropolitan parts of three states. Early used the resulting hedonic equations to predict median market rents of units with sample mean values of the regressors. The results reported in the last row of online table A-5 indicate that this price index is highly correlated with our basic index and the indices are very close to proportional on average. The mean absolute percentage difference is larger than in the previous comparisons, but still small.

⁷ In an earlier study, Moulton (1995) found that estimating separate regressions for different areas with the CPI sample led to poor out-of-sample predictions compared with a single regression that imposed the same coefficients on housing and neighborhood characteristics across areas. He attributed this result to small sample sizes in some areas.

In summary, the results reported in this appendix indicate that reasonable alternative methods for producing housing price indices with the CSS data yield indices that are very similar to our basic price index.

Appendix D

Approximations of Ceiling Rents Used to Estimate Truncated Regression Model

This appendix describes the approximations of the rent ceilings that faced the recipients of housing certificates and vouchers in our sample. These ceilings were used in the estimation of the truncated regression model discussed in appendix C.

Roughly 25 percent of the sample members were in the old certificate program. Most recipients under this program faced a rent ceiling equal to the local Fair Market Rent (FMR) for a unit with the number of bedrooms deemed appropriate for a family of its size and composition. Local housing authorities were allowed to approve rents up to 10 percent greater than the relevant FMR for up to 20 percent of recipients, and with HUD field office approval, they could allow rents up to 20 percent greater than the relevant FMR for these recipients. About 29 percent of certificate recipients were served by housing authorities that used exception rents at the time of our data [Devine et al., 2000, Table IV-8, Table A-2]. Recipients had a substantial incentive to find the best unit available renting for no more than their ceiling since occupying a more expensive unit, within that constraint, did not require them to sacrifice consumption of other goods. Our data contains information on the FMR that applied to each recipient, but not specific ceiling rents faced by recipients granted exceptions under the certificate program.

To approximate the preceding reality, we made the following assumptions. If the gross rent was less than or equal to the relevant FMR, the FMR was the ceiling rent. If the gross rent was greater than the FMR but less than or equal to $1.1 \cdot FMR$, the ceiling rent was $1.1 \cdot FMR$. If the gross rent was greater than $1.1 \cdot FMR$ but less than or equal to $1.2 \cdot FMR$, the ceiling rent was $1.2 \cdot FMR$. Finally, if the gross rent exceeded $1.2 \cdot FMR$, the ceiling rent was the gross rent.

The remaining 75 percent of sample members participated in the old or new voucher program. The CSS data does not distinguish between these programs. This distinction is important for our purposes because the old voucher program did not, and the new voucher program does, have a ceiling on the rent of the unit occupied. Based on other information, we conclude that about 10 percent of the CSS sample members were under the old voucher program and 65 percent under the new program. Since we could not determine which units were under the old program and these units were a distinct minority of all voucher units in the sample, we

assumed that all voucher units were covered by the rules of the new voucher program.

When a family enters the new voucher program or when it moves to a new unit under the program, it faces a ceiling on the rent of its unit equal to a local payment standard applicable to families of its type plus 10 percent of the family's adjusted income. Unlike the certificate program, these recipients do not have a strong incentive to occupy a dwelling unit renting for the ceiling rent. They bear the full marginal cost of more expensive housing for units renting for more than the local payment standard and less than the ceiling rent. Beyond the first year in a given unit, the rent can exceed this amount provided that the housing authority certifies that the rent does not exceed the market rent of similar units.

Since implementation of the 1998 Quality Housing and Work Responsibility Act, housing authorities have been allowed to establish local payment standards within 10 percent of the relevant FMR for some or all types of voucher recipients without HUD approval.⁸ Devine et al. (2000, p. 48) reports that 90 percent of housing authorities had adopted a uniform percentage of the FMR for all recipients at the time of our data.⁹ Among these housing authorities, about 64 percent used the FMR themselves as payment standards and about 21 percent used payment standards above FMR. The CSS data contains information on the FMR applicable to each recipient under the new voucher program, but it does not contain the applicable local payment standard. We approximated them based on a data file from HUD's Office of Public Housing and Voucher Programs that contains the payment standard applicable to each voucher recipient and other relevant information during our time period.¹⁰ It is not possible to match the households in this file with those in the CSS data. To approximate each housing authority's payment standard, we calculated separately for each housing authority the median payment standard among households living in the same zip code, with the same number of bedrooms on the voucher, and with and without a disabled member of the household. (Using the mean and mode payment standards produced similar results.) This allowed us to link an estimate of the payment standard at this level of specificity to more than 90 percent of voucher recipients in the CSS sample.

To approximate the ceiling rent for each household in the CSS data, we made the following assumptions. If the gross rent was less than or equal to $PS + .1 \cdot AINC$ (where $AINC$ is

⁸ With HUD approval, they could establish payment standards outside this range. However, few exceptions had been granted at the time of our data [Devine et al., 2000, p. 48].

⁹ Most of the rest established percentages that differed for families of different sizes and compositions and in different areas within their jurisdiction.

¹⁰ We are grateful to Milan Ozdinec and Juan Garcia for providing this information.

adjusted income), the ceiling rent was $PS + .1 \cdot AINC$. If the gross rent is greater than $PS + .1 \cdot AINC$, the ceiling rent was the gross rent. For the cases where the CSS did not report one of the variables needed to use the estimated payment standards for its locality (less than 10 percent of the cases), we assumed that the payment standard was $1.1 \cdot FMR$.

Appendix E

Construction of Price Indices for Other Goods and All Produced Goods

This appendix describes in detail how we create an interarea price index for all nonhousing goods and an overall consumer price index for all areas in 2000. Each quarter, ACCRA provides an overall cross-sectional consumer price index for many areas and price indices for most privately produced goods grouped into six categories.¹¹ However, its indices are not available for many other areas, and our housing price index is based on a much larger sample of dwelling units and is much better than the ACCRA index in accounting for differences in housing and neighborhood characteristics and avoiding errors in predicting the rental value of owner-occupied units. We use ACCRA nonhousing price indices, our housing price index, and other data to construct a price index for all nonhousing goods and an overall consumer price index for all locations in 2000.

The first step is to calculate an index of the price of all goods except housing and utilities for the places where the ACCRA index exists. To do it, we use ACCRA price indices for the four broad categories of other goods and average expenditure shares for all consumers from the Consumer Expenditure Survey (CEX). Online table A-6 reports our judgment about which CEX categories correspond to the four ACCRA nonhousing composite commodities based on an examination of the specific goods that ACCRA prices in each broad category. It also reports the expenditure share for each broad category used by ACCRA to create its overall consumer price index and the CEX expenditure share for all consumers. Our price index for nonhousing goods for the areas covered by ACCRA is the weighted mean of the ACCRA price indices for grocery items, transportation, health care, and miscellaneous goods using the CEX expenditure shares for all consumers as weights.

Our second step is to predict the price index for nonhousing goods for areas not covered by the ACCRA index. Our estimates are based on a simple theoretical model that recognizes that each good consumed in a locality involves some local labor and land and some imported inputs, often semi-finished or finished products. We assume that the production functions for

¹¹ Council for Community and Economic Research (2006) documents their data collection procedures and price index construction. It also sells a data file with the prices of the individual items used to create these price indices.

housing services H and other goods X are Cobb-Douglas with constant returns to scale, where output depends on the quantities of local labor L, local land K, imported inputs I, and inputs F whose prices are the same at all locations. Specifically,

$$Q_H = A_H Q_L^{\alpha_{LH}} Q_K^{\alpha_{KH}} Q_I^{\alpha_{IH}} Q_F^{\alpha_{FH}} \quad (1)$$

$$Q_X = A_X Q_L^{\alpha_{LX}} Q_K^{\alpha_{KX}} Q_I^{\alpha_{IX}} Q_F^{\alpha_{FX}} \quad (2)$$

where the A's and the α 's are constants.

These production functions imply the following minimum long-run average cost of production.

$$LRAC_H = (1/A_H)(P_L/\alpha_{LH})^{\alpha_{LH}} (P_K/\alpha_{KH})^{\alpha_{KH}} (P_I/\alpha_{IH})^{\alpha_{IH}} (P_F/\alpha_{FH})^{\alpha_{FH}} \quad (3)$$

$$LRAC_X = (1/A_X)(P_L/\alpha_{LX})^{\alpha_{LX}} (P_K/\alpha_{KX})^{\alpha_{KX}} (P_I/\alpha_{IX})^{\alpha_{IX}} (P_F/\alpha_{FX})^{\alpha_{FX}} \quad (4)$$

In the absence of government action, the long-run equilibrium prices of the two goods would be equal to these minimum long-run average costs.

Local government policies might affect output prices only through their effects on input prices. To account for the possibility that local government policies also create gaps between long-run equilibrium prices and minimum long-run average production cost at prevailing input prices, we assume that

$$P_H = E_H LRAC_H \quad (5)$$

$$P_X = E_X LRAC_X \quad (6)$$

where E_H and E_X are expected to be greater than or equal to 1.

Substituting (3) into (5) and (4) into (6) and taking the logarithm of both sides yields

$$\ln P_H = K_H + \ln E_H - \ln A_H + \alpha_{LH} \ln P_L + \alpha_{KH} \ln P_K + \alpha_{IH} \ln P_I + \alpha_{FH} \ln P_F \quad (7)$$

$$\ln P_X = K_X + \ln E_X - \ln A_X + \alpha_{LX} \ln P_L + \alpha_{KX} \ln P_K + \alpha_{IX} \ln P_I + \alpha_{FX} \ln P_F \quad (8)$$

where K_H and K_X are constants that depend on the α 's in the respective equations. Since P_F is the same everywhere, we can rewrite (7) and (8) as

$$\ln P_H = C_H + \ln E_H - \ln A_H + \alpha_{LH} \ln P_L + \alpha_{KH} \ln P_K + \alpha_{IH} \ln P_I \quad (9)$$

$$\ln P_X = C_X + \ln E_X - \ln A_X + \alpha_{LX} \ln P_L + \alpha_{KX} \ln P_K + \alpha_{IX} \ln P_I \quad (10)$$

where C_H and C_X are constants.

If data were available on the three composite input prices and determinants of E_X and A_X , it would be possible to estimate (10) using data for locations where P_X is reported and use this estimated regression equation to predict this variable for other locations.¹² Equation (9) would be irrelevant. However, data on land prices P_K and the prices of imported inputs P_I are not readily available. To account for these unobserved input prices, we first solve (9) for P_K and substitute into (10). This yields

$$\begin{aligned} \ln P_X = & [C_X - (\alpha_{KX} / \alpha_{KH}) C_H] + [\ln E_X - (\alpha_{KX} / \alpha_{KH}) \ln E_H] + [(\alpha_{KX} / \alpha_{KH}) \ln A_H - \ln A_X] \\ & + [(\alpha_{LX} \alpha_{KH} - \alpha_{LH} \alpha_{KX}) / \alpha_{KH}] \ln P_L + (\alpha_{KX} / \alpha_{KH}) \ln P_H + [(\alpha_{IX} \alpha_{KH} - \alpha_{IH} \alpha_{KX}) / \alpha_{KH}] \ln P_I \end{aligned} \quad (11)$$

With the exception of the second term, the terms in square brackets reflect differences in the parameters of the production functions for housing services and other goods. If local government policies affected output prices only through their effect on input prices ($E_H = E_X = 1$) and there were no differences in production functions, these terms would be zero, the coefficient of $\ln P_H$ would be 1, and the prices of the two goods would be the same in each location. In this case, our housing price index would also be a price index of nonhousing goods

¹² The error term in this regression model stems from error terms in equations explaining $\ln E_H$ and $\ln A_H$.

and all goods. If there are differences in production functions for the two goods, the inclusion of the price of housing services in a regression model explaining differences in the price of other goods is useful because it captures the effect of unobserved input prices, especially land prices.

To complete the regression model, we first write $\ln E_H$ and $\ln E_X$ as functions of an index of land use regulation (*regindex*). Regulations might create a deviation between price and production cost due, for example, to expenditures to secure variances from these regulations. We write $\ln A_H$ and $\ln A_X$ as functions of climate variables (*coolingdays*, *heatingdays*, *precip*) because weather might affect the output that can be produced with a given input bundle. Finally, we write $\ln P_I$ as a function of the distance to the nearest metropolitan area with a population in excess of 1.5 million (*dist*). Areas that are farther from large metropolitan areas would arguably have higher prices of imported inputs. Each equation is assumed to have an additive error term with standard properties. Appendix F provides the definitions and sources of these variables. Substituting these equations into (11) and reparameterizing yields the regression model used to explain differences in the price index for other goods across areas where it was available.

$$\begin{aligned} \ln P_X = & \beta_0 + \beta_1 \text{regindex} + \beta_2 \ln(\text{coolingdays} + 1) + \beta_3 \ln(\text{heatingdays} + 1) \\ & + \beta_4 \text{precip} + \beta_5 \text{precip}^2 + \beta_6 \ln P_L + \beta_7 \ln P_H + \beta_8 \text{dist} + \varepsilon \end{aligned} \quad (12)$$

Table 3 reports the OLS estimates of the parameters of this model.¹³ Under plausible assumptions about the underlying error terms, the error term in (14) would be correlated with $\ln P_H$ and hence OLS estimators of the β coefficients would be biased. However, since the purpose of this estimated equation is to predict the index of nonhousing prices where it is not reported based on data available, this is not problem with OLS estimation.

Analysis of the residuals suggests no significant misspecification of functional form, heteroskedasticity, or outliers. If the functional form is correct, we expect the mean of the residuals to be about zero at all predicted values of $\ln P_X$. Across the quintiles of the distribution of these predicted values, the mean of the residuals ranged from -.0141 to .0086. If the error term in the regression model is homoskedastic, the residuals should have about the same

¹³ Because the ACCRA prices for New York City refer to Manhattan, an unusually expensive part of the NYC PMSA, we treat these prices as not reported in estimating the model and predict the nonhousing price index for it.

standard deviation at all predicted values of $\ln P_X$. Across the quintiles of these predicted values, the standard deviations of the residuals ranged from .0328 to .0406. The largest deviation between predicted and observed value of the price index for other goods is 10 percent and less than a quarter exceed 5 percent.

It is important to realize that the insignificance of the regulation index in our regression model does not indicate that these regulations have no effect on nonhousing prices. The relevant regression model for this purpose would exclude $\ln P_H$ from (12). It stems from writing $\ln P_K$ in (10) as a function of *regindex*, *coolingdays*, *heatingdays*, and *precip* and $\ln P_I$ as a function of *dist*. Similar substitutions into (9) yield the equation relevant for estimating the effect of regulations on housing prices. The coefficient of *regindex* in these regressions capture the effects of land use regulations on output price that operate through their effects on land prices as well as any gaps that they create between long-run equilibrium prices and minimum long-run average production cost at prevailing input prices. The OLS estimate of the coefficient of *regindex* in the regression model explaining nonhousing prices is .01 with a t-score of 2.06. The effect of land use regulations on housing prices is much greater; the estimated coefficient is .06 with a t-score of 4.86. Assuming no effect of these regulations on wage rates, these results imply that a one standard deviation increase in the regulation index results in a one percent increase in the price of nonhousing goods and a six percent increase in the price of housing services. The substantial effect of land use regulations on housing prices is consistent with the bulk of the previous literature, for example, Glaeser and Gyourko (2003), Glaeser, Gyourko, and Saks (2005), Glaeser and Ward (2009), Katz and Rosen (1987), Malpezzi (1996), and Quigley and Raphael (2005).

Table 1 reports the price indices for housing, other goods, and all goods in 2000 for the ten areas with the highest, lowest, and middle housing price index. Each price index is scaled to have a mean of 1 across all locations. The price index for nonhousing goods is the rescaled ACCRA index for the areas where it was available and the predicted index for other areas. The overall price index is the weighted average of the price indices for housing and other goods where the weights are the CEX expenditure shares for all consumers. Online table A-4 reports these price indices for all locations. Table 1 suggests what is generally true. On average, nonhousing prices are higher in areas where housing prices are higher, and the ratio of housing prices to the prices of nonhousing goods are higher in areas with the highest overall CPI. The

highest housing price index is three times as large as the smallest. The highest price index for other goods is only 39 percent greater than the smallest.

Since some researchers will want to use the overall consumer price index to study subsets of the population, it is worthwhile to determine its sensitivity to the weights used to construct it. The ACCRA price indices are based on expenditure weights that reflect the consumption patterns of a very special subset of the population. One reason that economists have been reluctant to use the ACCRA index is that they are studying different populations with different expenditure patterns and they believe that price indices would be sensitive to these differences. As mentioned earlier, Koo, Phillips, and Sigalla (2000, pp. 130-131) have found that replacing ACCRA's expenditure weights with weights reflecting average expenditure shares has very little effect on the overall price index, albeit in a study limited to 23 metropolitan areas. The results of our study based on 380 areas supports their conclusion. When we compare an overall price index using the ACCRA expenditure shares in online table A-6 with our price index based on the very different expenditure shares of all consumers from the CEX, the resulting indices are virtually identical. The correlation coefficient between the two price indices exceeds .99, the largest percentage difference between the two is less than 7 percent, and the mean absolute percentage difference is less than 2 percent.

The simple formula used to calculate our overall price index is not ideal from the viewpoint of measuring differences or changes in well-being. The ratio of an individual's income to this price index is not an index of the individual's well-being for any preferences with the standard general properties [Deaton and Muellbauer, 1980, Chapter 7; Pollak, 1989, Chapter 1].¹⁴ In a simple world in which income is not subject to choice and there are no differences in amenities across locations and individuals face budget frontiers that are hyperplanes, an ideal price index for an individual exists if and only if the individual's indirect utility function can be written as the ratio of the individual's income to an expression involving only the prices of goods and constants. The expression in the denominator is an ideal price index. If there are differences in amenities across locations and amenities are separable from other goods in the person's utility function, then the ratio of income to the price index is an index of the well-being that results from consumption of the goods priced. An index of overall well-being would require accounting for differences in amenities across locations.

¹⁴ Furthermore, when it exists at all, an ideal price index is different for each person. We ignore this complication.

To get some sense of whether moving towards an ideal price index would yield very different results, we develop an ideal overall price index based on a simple assumption about preferences and compare it with our price index. The ideal price index is based on the assumption that all people have a Cobb-Douglas utility function involving two goods housing and nonhousing with exponents equal to the expenditure shares that underlie the previous overall price index. The formula for this price index is:

$$CPI = (PH / .252)^{.252} (PX / .748)^{.748} \quad (3)$$

After rescaling this ideal price index to have the same mean as the simple expenditure weighted average of the housing price index PH and the price index of other goods PX in online table A-4, the price indices are almost identical. The correlation coefficient exceeds .999, the mean absolute percentage difference is less than three-tenths of a percent, and the maximum absolute percentage difference is 2.3 across the 380 locations.

Appendix F

Explanatory Variables in Regression Explaining Differences in Nonhousing Prices

This appendix documents the sources of the explanatory variables in the regression explaining differences in the nonhousing price index and how values were imputed when they were not reported.

Land use regulation index (regindex)

We estimated a regulatory index for our areas using the Wharton Residential Land Use Regulatory Index (WRI) developed by Gyourko, Saiz and Summers (2007). This index is based on a nationwide survey of local land use controls. The survey was sent to about 6,900 municipalities across the U.S. About 38 percent responded, representing about 60 percent of the surveyed population. The survey data together with information on state land use policies and other measures of community pressure (using information from environmental and open space ballot initiatives) are used to create eleven subindexes that summarize different aspects of land use regulation. Higher values of these indices indicate more restrictive regulations. An aggregate index, the WRI, is created using factor analysis. The WRI is standardized so that its sample mean is zero and standard deviation equals one. We use their municipal-level WRI index and weights (which are available online) to compute (weighted) average regulatory indices for most of our areas. Forty of our areas contain no subareas for which the WRI is reported. To impute the regulation index for 37 of these areas, we use a state-level average WRI provided by the authors in the paper. For three areas whose boundaries cross state lines (Cumberland, MD-WV MSA, Grand Forks, ND-MN MSA, and Clarksville-Hopkinsville, TN-KY MSA), we use simple averages of the corresponding state level regulation indices.

Climate (coolingdays, heatingdays, precip)

Census Bureau (2007, Table C-6) provides average annual precipitation and the total number of cooling and heating days between 1970 and 2000 for many cities. The level of geography in our study is the metropolitan area (MSA or PMSA, hereafter MSA), and the non-metropolitan part of each state. Metropolitan areas often contain multiple cities, but MSA names usually include the

name of its largest city. For these MSA, our values of the climate variables are the values for the largest city. In 18 cases, the source did not contain data for the cities mentioned in the MSA name or the MSA name contains only counties. In these cases, our imputed values of the climate variables were for the closest MSA. The median distance from the center of these 18 MSA to the closest MSA whose climate data were reported was 25 miles; the maximum was only 53 miles. The imputed values for the non-metropolitan part of each state are the mean values of the variables for the MSA in the state.

Wage rate (P_L)

Using U.S. Census data from the 2000 Integrated Public-Use Microdata Series (IPUMS), David Albouy (2009) computed wage differentials across 290 areas of the U.S. The wage differentials are computed for full time workers (working at least 30 hours a week, 26 weeks a year) ages 25 to 55. To estimate the wage differential, a log-wage regression is estimated. Covariates include educational attainment, potential experience, industry, gender, English proficiency, and marital, veteran, minority, and immigrant status, their interactions, and MSA dummy variables. The regression model is estimated using weighted OLS. Albouy's index is an index of $\ln P_L$, whose value is zero in Reading, PA.

Some assumptions are necessary to predict wage indices for all of our areas. Albouy computes only one wage index for each Consolidated Metropolitan Statistical Area (CMSA). We assume that all PMSAs within the CMSA have the same wage index. Wage data are unavailable for 35 small MSAs. The average population in these areas is about 120,000; the largest is Huntington-Ashland with about 315,000 residents. We use the wage index computed for the non-metropolitan part of the corresponding state to impute the wage index of these small areas. In the eight cases where the MSA spans several states, we compute a simple average of the corresponding non-MSA state wage indices.

Distance to nearest large metropolitan area ($dist$)

This variable is the 'as-the-crow-flies' distance between the center of each area and the center of the nearest MSA with at least 1.5 million residents. For the 41 large metropolitan areas, it is zero. The center of the non-metropolitan part of each state is assumed to be the center of the state. The longitude and latitude of the center of each area were obtained from Google Maps.

Appendix References

- Albouy, David. 2009. "The Unequal Geographic Burden of Federal Taxation." *Journal of Political Economy* 117(4): 635-667.
- Albouy, David. 2012. "Are Big Cities Really Bad Places to Live? Improving Quality-of-life Estimates across Cities." Unpublished manuscript.
- Aten, Bettina H., Eric B. Figueroa, and Troy M. Martin. 2011. "Regional Price Parities by Expenditure Class, 2005-2009." *Survey of Current Business* 91(5): 73-87.
- Baum-Snow, Nathaniel, and Ronni Pavan. 2012. "Understanding the City Size Wage Gap." *Review of Economic Studies* 79(1):88-127.
- Blackley, Dixie M., and James R. Follain, with Haeduck Lee. 1986. "An Evaluation of Hedonic Price Indexes for 34 Large SMSAs." *AREUEA Journal* 14(2): 179-205.
- Building Research Council, School of Architecture, University of Illinois at Urbana-Champaign. 1998. *Resident Assessment of Housing Quality: Lessons from Pilot Surveys*. Washington, DC: U.S. Department of Housing and Urban Development.
- Chen, Yong, and Stuart S. Rosenthal. 2008. "Local Amenities and Life-Cycle Migration: Do People Move for Jobs or Fun?" *Journal of Urban Economics* 64(3): 519-537.
- Council for Community and Economic Research. 2006. *ACCRA Cost of Living Index Manual*. Arlington, VA:
- Deaton, Angus, and John Muellbauer. 1980. *Economics and Consumer Behavior*. Cambridge: Cambridge University Press, 1980.
- Devine, Deborah J., Barbara A. Haley, Lester Rubin, and Robert W. Gray. 2000. *The Uses of Discretionary Authority in the Tenant-Based Section 8 Program*. Washington, DC: U.S. Department of Housing and Urban Development.
- Early, Dirk W. 2006. "Metropolitan Area Cost of Rental Housing Indices Using the Customer Satisfaction Survey," HUD Report, Research for Assisted Housing II Grants Program, Grant Number GS-10F-0405M, US Department of Housing and Urban Development, Washington, D.C.
- Follain, James R. 1979. *How Well Do Section 8 FMRs Match the Cost of Rental Housing?: Data from 39 Large Cities*. Contract Report 249-15. Washington, D.C.: The Urban Institute.
- Follain, James R., and Larry Ozanne, with Verna Alberger. 1979. *Place to Place Indexes of the*

- Price of Housing*. Washington, DC: The Urban Institute.
- Gabriel, Stuart, and Stuart S. Rosenthal. 2004. "Quality of the Business Environment versus Quality of Life: Do Firms and Households Like the Same Cities?" *Review of Economics and Statistics* 86(1): 438-444.
- Glaeser, Edward L., and Joseph Gyourko. 2003. "The Impact of Building Restrictions on Housing Affordability." *Economic Policy Review* 9(2): 21-39.
- Glaeser, Edward L., Joseph Gyourko, Raven A Saks. 2005. "Why Is Manhattan So Expensive? Regulation and the Rise in Housing Prices." *Journal of Law and Economics* 48(2): 331-369.
- Glaeser, Edward L., and Bryce A. Ward. 2009. "The Causes and Consequences of Land Use Regulation: Evidence from Greater Boston." *Journal of Urban Economics* 65(3): 265-278.
- Gray, Robert W., Barbara A. Haley, Seth R. Marcus, Huda M. Elhassan, and Vivienne Liu. 2002. *Tell Us About Your Home: First Year Results of the Survey of Households Assisted by the Housing Choice Voucher Program, Draft Report*. Program Monitoring and Research Division, Office of Policy Development and Research, Washington, DC: U.S. Department of Housing and Urban Development.
- Gyourko, Joseph, Albert Saiz and Anita Summers. 2008. "A New Measure of the Local Regulatory Environment for Housing Markets: The Wharton Residential Land Use Regulatory Index." *Urban Studies* 45(3): 693-729.
- Johnson, David S., John M. Rogers, and Lucilla Tan. 2001. "A Century of Family Budgets in the United States." *Monthly Labor Review* 124(5): 28-45.
- Katz, Lawrence and Kenneth T. Rosen. 1987. "The Interjurisdictional Effects of Growth Controls on Housing Prices." *Journal of Law and Economics* 30(1): 149-160.
- Kiel, Katherine, and Jeffrey E. Zabel. 1997. "Evaluating the Usefulness of the American Housing Survey for Creating House Price Indices." *Journal of Real Estate Finance and Economics* 14: 189-202.
- Kokoski, Mary F., Patrick Cardiff, and Brent R. Moulton. 1994. "Interarea Price Indices for Consumer Goods and Services: A Hedonic Approach Using CPI Data." U.S. Department of Labor, Bureau of Labor Statistics, Working Paper 256.
- Koo, Jahyeong, Keith R. Phillips, and Fiona D. Sigalla. 2000. "Measuring Regional Cost of

- Living.” *Journal of Business & Economic Statistics*, 18(1): 127-136.
- Maddala, G. S. 1983. *Limited-Dependent and Qualitative Variables in Econometrics*.
Econometric Society Monographs, Cambridge, UK: Cambridge University Press.
- Malpezzi, Stephen. 1996. Housing Prices, Externalities, and Regulation in U.S. Metropolitan Areas.” *Journal of Housing Research* 7(2): 209-241.
- Malpezzi, Stephen, Gregory Chun and Richard Green. 1998. “New Place to Place Housing Price Indexes for U.S. Metropolitan Areas, and Their Determinants: An Application of Housing Indicators,” *Real Estate Economics* 26(2): 235-75.
- Moretti, Enrico. 2013. “Real Wage Inequality.” *American Economic Journal: Applied Economics* 5(1): 65-103.
- Moulton, Brent R. 1995. “Interarea Indexes of the Cost of Shelter Using Hedonic Quality Adjustment Techniques,” *Journal of Econometrics* 68(1): 181-204.
- Olsen, Edgar. 2003. “Housing Programs for Low-Income Households.” In *Means-Tested Transfer Programs in the United States*, ed., Robert Moffitt. Chicago: University of Chicago Press.
- Pollak, Robert A. 1989. *The Theory of the Cost-of-Living Index*. Oxford: Oxford University Press.
- Ptacek, Frank, and Robert M. Baskin. 1996. “Revision of the CPI Housing Sample and Estimators.” *Monthly Labor Review* 119(12): 31-39.
- Quigley, John M. and Steven Raphael. 2005. “Regulation and the High Cost of Housing in California.” *American Economic Review* 95(2): 323-328.
- Thibodeau, Thomas G. 1989. “Housing Price Indexes from the 1974-83 SMSA Annual Housing Surveys.” *AREUEA Journal* 17(1): 100-17.
- Thibodeau, Thomas G. 1995. “House Price Indices from the 1984-1992 MSA American Housing Surveys,” *Journal of Housing Research* 6(3): 439-479.
- U.S. Census Bureau. 2007. *County and City Data Book: 2007*. 14th edition. Washington, D.C.
- Winters, John V. 2009. “Wages and Prices: Are Workers Fully Compensated for Cost of Living Differences?” *Regional Science and Urban Economics* 39(5): 632-643.