Improving Median Housing Price Indexes through Stratification

Authors   Nalini Prasad and Anthony Richards

Abstract   There is a trade-off between how easy a housing price series is to construct and the extent to which it adjusts for changes in the mix of dwellings sold. Median house price measures are easily calculated, frequently used by industry bodies, and quoted in the media. However, such measures provide poor estimates of short-term changes in prices because they reflect changes in the composition of transactions, as well as changes in demand and supply conditions. This study uses a database of 3.5 million transactions in the six largest Australian cities to demonstrate that compositional shifts between higher- and lower-priced parts of cities can account for much of the noise in median price measures. Accordingly, a simple method of adjusting for compositional change through stratification is proposed. The measure differs from those commonly used internationally, as neighborhoods or small geographic regions are grouped according to the long-term average price level of dwellings in those regions. The measure of price growth produced improves substantially upon a median and is very highly correlated with regression-based measures.

Developments in housing prices are of intense interest to households, policymakers, and those involved in the housing industry. This has been particularly true recently, as rapid increases in housing prices across a number of countries have led to concerns about housing markets being overvalued and to fears that this might be followed by a sharp correction in prices. In such circumstances, it becomes very important to have good and timely measures of short-term movements in housing prices. One particular reason is that macroeconomic policy settings that were appropriate when housing prices were rising rapidly could suddenly be extremely inappropriate if the market was now falling.[1]

However, the construction of aggregate measures of housing prices is not a straightforward exercise. One major problem in measuring housing price growth results from the infrequency of transactions and the heterogeneous nature of the housing stock. Only a relatively small fraction of the housing stock is transacted in any period. For example, in the United States and Australia the average turnover

for the existing dwelling stock is around 5% per year, or just 1.25% per quarter and in other countries the turnover rate is often significantly lower. Given that the sample of transactions in any period may not be representative of the entire housing stock, it becomes important to ensure that measures of house prices reflect true movements in the housing market rather than spurious movements due to compositional effects.

One measure of prices that is widely calculated is the median (or mean) of transaction prices. For example, the U.S. National Association of Realtors, the Canadian Real Estate Organization, the Real Estate Institute of Australia, and the Real Estate Institute of New Zealand all publish house price data that are simple median or mean measures. However, conventional wisdom among researchers would most likely dismiss median price measures as good measures of housing prices.[2] Although median price measures typically score highly on timeliness, they tend to be significantly affected by compositional change. Instead, the conventional wisdom would no doubt suggest using regression-based approaches such as hedonics or repeat-sales analysis to abstract from compositional effects and derive estimates of pure price changes.[3] However, the question arises as to whether these latter measures can perform as well in real-time situations as they can in more ideal circumstances. For example, initial estimates from repeat-sales regressions can be subject to significant revision, because initial estimates of price growth for any quarter will be based on only a fraction of the sales that will eventually influence the estimate.[4] And hedonic price regressions may not be feasible in real-time if data on transactions prices cannot be immediately matched with data on property characteristics. Given the data requirements of these more sophisticated approaches, it is not surprising that such measures tend to become available on a less timely basis than median measures.[5]

Given these potential real-time difficulties with the more sophisticated methodologies, it becomes worthwhile to consider if median price measures can be improved to be more useful before reliable estimates can be produced using the more sophisticated regression techniques. This paper takes up that task, drawing on the significant evidence (e.g., Straszheim, 1975) that location is a major determinant—if not the most important single determinant—of housing prices. Since median price measures do not control for the location of a dwelling within a city, and since there can be large differences in prices across different parts of cities, it seems plausible to conjecture that locational effects could be responsible for much of the compositional effects that cause simple median price measures to yield poor estimates of short-term price movements. If so, it might be possible to control for such locational effects to derive measures based on median prices that yield estimates of short-term price movements that are good, timely, and easy to compute.

This study uses data for house prices in the six largest Australian cities and apartment prices in the two largest cities. The dataset used contains approximately 3.5 million transactions over 1993–2005. The findings confirm that simple city-wide medians provide poor estimates of short-term price changes and that this is

substantially due to compositional change in the mix of sales between higher- and lower-price neighborhoods within these cities. A simple method for calculating changes in aggregate city-wide housing prices that explicitly controls for compositional change is tested. The particular innovation of the paper is the method of stratification that is used. The method divides a city into broad geographical regions. However, changes in regional composition do not necessarily result in problems for median measures; compositional change will only be a significant problem if it results in changes in the proportion of high- and low-priced properties. Accordingly, small geographical regions within cities are grouped into different strata based on the long-term average price level of houses (or apartments) in those regions, thereby directly addressing the main problem of compositional change. Stratifying sales in this manner produces a measure of price growth that is a considerable improvement over an unstratified median; the measure is significantly less noisy than a median and performs better with limited data samples (that is, the ones available to policy makers in real time). The findings also reveal that the growth rates produced by this methodology line up closely with estimates based on hedonic and repeat-sales approaches. The advantage of this approach is that it is easy to compute because it is based on simple medians from stratification and uses data—on transactions price and location—that are readily available from most housing transactions databases. In summary, the paper demonstrates that it is possible to generate good estimates of short-term price movements from median prices, if the medians are taken from an appropriately stratified data sample that is designed to address the key problems of compositional change.

The rest of the paper is organized as follows. An overview is provided of the Australian housing price data used in the paper and then the importance of location in determining housing prices is discussed, along with the impact of compositional change on Australian median housing price data. Following this, the method of controlling for compositional change is outlined, along with an assessment of the resulting measure of housing prices.

Data

This paper constructs a measure of housing price growth for the six largest Australian capital cities (Sydney, Melbourne, Brisbane, Adelaide, Perth, and Canberra—in descending order of population) covering the period from 1993:Q1 to 2005:Q4.[6] Sales information for this study is supplied by a private data provider, Australian Property Monitors (APM), which sources information from official government databases. In addition, since there are reporting lags for the data for the most recent quarters, APM supplements those quarters with its own data collection from real estate agents (or realtors).

The full dataset contains approximately 3.5 million observations, which is virtually the entire population of transactions that took place during the study period. The data for Sydney and Melbourne include sales of both houses (comprised

of detached and semi-detached dwellings) and apartments (which, in U.S. terminology, include both condominiums and co-ops). The stock of apartments in the other cities is much smaller, so for those cities the analysis in the paper covers only sales of houses.

The six cities are all defined to include the broader metropolitan area rather than just the inner city. As of June 2005, the populations of the six cities ranged from approximately 4.25 million (Sydney) to 300,000 (Canberra). Average sales volumes range from around 15,800 house sales per quarter in Melbourne to around 1,200 house sales per quarter in Canberra. These sample sizes are quite large, reflecting the fact that turnover in the Australian housing market is relatively high by international standards, yet as shown below, city-wide median prices are still noisy at a quarterly frequency.

The data within each city can be split on the basis of the 'suburb' (or small neighborhood in U.S. terminology) where a property is located. This is the finest level of geographical disaggregation available (and is a finer level of disaggregation than postcodes/ZIP Codes). The number of suburbs ranges from 659 in Sydney to 84 for Canberra. On average, there are around 5,500 people and around 2,000 dwellings per suburb, indicating a fairly fine level of disaggregation is possible within each city.

The Impact of Compositional Change on Median Price Measures

Median or mean house price series are produced in many countries. One clear advantage of median price measures is that they are very easy to calculate. They also have a straightforward interpretation: they represent the price of a 'typical' transaction in any given period. However, if one is interested in inferring the price change for the overall housing stock, these measures can be distorted by compositional change and the extreme range of observed housing prices. Transactions that occur in any period may not be representative of the overall housing stock. Importantly for estimating price changes, the composition of the sample of transactions in one period may be quite different to the composition in the next period.

The Importance of Location in Determining Prices

Just as realtors stress the importance of 'location, location, location,' there is a general (and growing) consensus in the academic literature that a key characteristic of urban housing markets is the variation of prices by location (e.g., Straszheim 1975; and Goodman and Thibodeau, 2003). As is discussed further by Can (1998), the geographic location of a dwelling determines households' access to all sorts of amenities and services, so there can be sizeable differences in the prices of houses of a similar physical condition and structure depending on their location.

Individuals are often willing to pay a premium for desirable locations. Not surprisingly, a large proportion of the explanatory power in hedonic regressions typically comes from variables describing the location of each dwelling.

Various authors, including Can and Megbolugbe (1997) and Berg (2005) have shown that the fit of hedonic models is much improved when information on recent selling prices of nearby properties is added. However, while there are various sophisticated spatial techniques for identifying and grouping together homogenous transactions, a number of authors have found that standard geographical boundaries provide a valid alternative. For example, Bourassa, Cantoni, and Hoesli (2005) find that submarkets defined according to geography are better for predicting house prices than those formed using more complicated spatial statistical models. Similarly Goodman and Thibodeau (2003) find that defining submarkets by ZIP Codes produces comparable results (in terms of predicting house prices) to using hierarchical models to define submarket boundaries.

An examination of the Australian data on transactions prices was conducted to identify some stylized facts about the role of location in influencing dwelling prices in different cities. The data are disaggregated into 'suburbs' or neighborhoods, but as even the smallest city in the sample has over 80 suburbs, these small geographic regions are grouped in such a way as to facilitate the analysis. In particular, suburbs within each market were ranked according to their median transactions price over 2000–2004. This ordering of median prices was then used to group suburbs into deciles, with an equal number of suburbs in each decile. For example, of the 659 suburbs in Sydney, Decile 1 contains the transactions of the 65 suburbs with the lowest median house prices and Decile 10 consists of the transactions in the 66 suburbs with the highest median prices. Deciles are formed in each city market, with the exception of the house market in Canberra and the apartment markets in Sydney and Melbourne. Given the smaller samples of transactions in these markets, transactions were grouped into quintiles using the same criteria outlined above. Each decile (or quintile) now contains properties from regions that are similar in terms of median price.

Exhibit 1 shows the median price (in Australian dollars with a log scale) for each decile in the largest five housing markets for 2005, which is the year immediately following the ranking period. There is a wide range of variation in the prices within each city. The median price of Decile 10 in Sydney is around 4.5 times higher than prices in the bottom decile. The range in the other cities is also notable, although not as large, with median prices in the top decile typically around 3.2 times larger than in the bottom decile. For the smaller markets (not shown in the graph), median prices in the top quintile are roughly double the price of those in the lower quintile.

Given that city-wide median (or mean) price measures do not control for the location of properties that are transacted, they can be influenced by changes in the locational composition of the transactions that occur in any period. In

**Exhibit 1** | Distribution of Housing Prices in Australian Cities
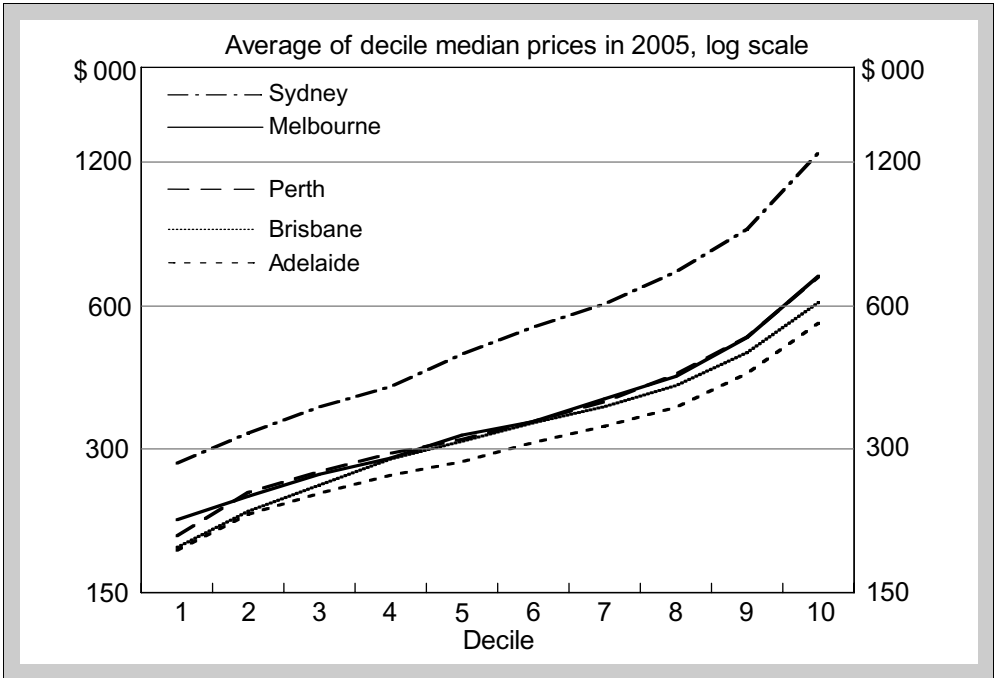
Source: Author's analysis using data from APM.

particular, as demonstrated above, cities have areas where dwellings tend to be more expensive and other areas where they tend to be less expensive. If the mix of transactions between these groups varies significantly, it would be expected to impact median price measures. As a result, changes in median and mean prices may contain substantial noise from locational composition change and provide poor estimates of true price changes.

Australian Evidence on Compositional Effects on Median Prices

The problems resulting from compositional change are illustrated using quarterly data for the city-wide median of prices of houses transacted in Sydney. The top panel of Exhibit 2 shows the quarterly median price (on a log scale) over 1993–2005. The middle panel shows the quarterly change in this series along with a line for the trend quarterly growth.[7,8] There has been substantial growth in median prices over most of this period, but with substantial noise, which is apparent in the saw-tooth pattern and the frequent large divergences between the actual and trend change in the median.

Exhibit 2 | Sydney Median House Prices

The bottom panel of Exhibit 2 includes a measure of compositional change, which may be able to explain some of the noise in the median price. The variable was constructed by calculating the proportion of transactions that took place in more expensive suburbs (Deciles 6–10). In the case of Sydney, this proportion averages somewhat below 50% because the allocation of suburbs was done to ensure a similar number of suburbs, rather than transactions, in each decile. The data show that there is significant quarterly variation in the proportion of transactions in the higher-priced suburbs.[9] Hence growth in the city-wide median price will reflect changes that result from compositional effects as well as pure price changes.

The proposition that compositional change between higher- and lower-priced suburbs may be responsible for some of the observed noise in the change in

median house prices can be examined. Quarterly changes in median prices are regressed on a constant and a compositional change variable, given by the quarterly change in the proportion of transactions that were in the more expensive suburbs. This equation is estimated for median housing prices in each city market. In most cases, the proportion of house sales in more expensive suburbs is based on sales in the top five deciles. In the cases where transactions were grouped into quintiles, the middle quintile is classified in such a way as to have an approximately equal number of sales in the higher- and lower-priced segments.

Results are shown in Exhibit 3. In all cases, the compositional change variable takes the expected positive sign: an increase in the proportion of transactions in higher-priced suburbs leads to the change in the city-wide median price being higher.[10] For Sydney and Melbourne, the results indicate that a considerable proportion—around 60%—of the quarterly variation in the city-wide median house price can be explained purely by shifts in the mix of sales between higher- and lower-priced suburbs. The effect of compositional change is less pronounced in the other markets, but (with the exception of Brisbane) it still explains around 20% to 40% of quarterly price movements. Indeed, on a national level, the

Exhibit 3 | Testing for the Impact of Compositional Change on Median Prices

	Coeff. on Compositional Change Variable	Adj. R^2
Sydney houses	1.09***	0.60
Melbourne houses	0.77***	0.63
Brisbane houses	0.56*	0.05
Perth houses	0.70***	0.20
Adelaide houses	0.89***	0.26
Canberra houses	0.45***	0.18
Sydney apartments	0.70***	0.37
Melbourne apartments	0.92***	0.39
Australian housing	1.03***	0.56

Notes: This table shows the results from a regression to determine if the quarterly growth in median house prices over 1993:Q2–2005:Q3 is affected by changes in the composition of dwellings sold. The regression estimated is $\Delta p_t = \alpha + \beta \Delta com_t + \varepsilon_t$, where Δp_t is the quarterly change in median prices, Δcom_t is the quarterly change in the proportion of transactions in more expensive suburbs (or neighborhoods), and εt is the error term.
*Significant at the 1% level.
**Significant at the 5% level.
***Significant at the 10% level.

compositional change variable can explain 56% of the quarterly movements in nationwide median prices, where these are calculated as a weighted average of prices using dwelling stock weights. Overall, the results in Exhibit 3 suggest fairly strongly that there may be significant gains from taking account of the effect of this simple form of compositional change on median price measures.

Part of the quarterly variation in the composition of transactions may be seasonal in nature. Accordingly, seasonality is tested for in the composition of sales and in the level of prices by using the X12 seasonal adjustment program.[11] In addition, the quarterly change in the proportion of transactions in more expensive suburbs is regressed, along with the quarterly change in the median price on seasonal dummies, and the adjusted R^2 from these regressions is used as a shorthand way to compare the importance of deterministic seasonality in different series. Panel A of Exhibit 4 contains results from testing for seasonality in the proportion of houses sold in more expensive suburbs and Panel B contains the results for median prices.

Both median prices and the composition of transactions are found to be seasonal in most markets. Furthermore, the markets where the composition of transactions is found to exhibit a seasonal pattern tend to be those where median prices are found to be seasonal, suggesting that at least part of the seasonality seen in median prices is the result of seasonality in the composition of sales.[12] The signs of the seasonal factors on the two variables (not shown, but available upon request) also support this, with the quarters when median prices are seasonally high tending to be the quarters when the proportion of sales in higher-priced suburbs is also seasonally high. The values for the adjusted R^2s suggest that seasonal influences are particularly strong in the market for houses in the two largest cities, explaining as much as a third (Sydney) or half (Melbourne) of the variation in quarterly price movements. Given that there is significant seasonality in most cities, and that the pattern of seasonality in the two largest cities is very similar, it is not surprising that there is also seasonality in average nationwide prices. Seasonal factors can explain nearly 40% of quarterly median price movements at the national level.

However, the relationship between changes in median prices and changes in the proportion of houses sold in more expensive suburbs is not purely due to common seasonality. For each city market, quarterly changes in seasonally adjusted median prices have been regresses on a constant and the seasonally adjusted compositional change variable (with results available upon request). The adjusted R^2s from these regressions are lower than those in Exhibit 3. However, in nearly all cities (the exceptions are Brisbane and Adelaide), the seasonally adjusted compositional change variable can explain a notable amount of the quarterly change in the seasonally adjusted median price, with adjusted R^2s ranging from between 0.10 (for Perth houses) to 0.37 (for Melbourne apartments). Therefore, there also exist significant non-seasonal shifts in the proportion of sales in more and less expensive suburbs that are reflected in movements in city-wide median prices.

Exhibit 4 | Testing for Seasonality

	Is Seasonality Present in X12?	Adj. R^{2a}
Panel A: Testing for seasonality in the compositional change variable		
Sydney houses	Yes	0.60
Melbourne houses	Yes	0.83
Brisbane houses	No	0.14
Perth houses	Yes	0.29
Adelaide houses	Yes	0.47
Canberra houses	No	0.05
Sydney apartments	Yes	0.53
Melbourne apartments	Yes	0.31
Australian housing	Yes	0.78
Panel B: Testing for seasonality in median house prices		
Sydney houses	Yes	0.33
Melbourne houses	Yes	0.50
Brisbane houses	No	0.02
Perth houses	Yes	0.14
Adelaide houses	Yes	0.26
Canberra houses	No	−0.02
Sydney apartments	Yes	0.26
Melbourne apartments	No	0.09
Australian housing	Yes	0.39
Panel C: Testing for seasonality in selected international median and mean price series		
U.S.		
Northeast	Yes	0.23
Midwest	Yes	0.45
South	Yes	0.55
West	Yes	0.18
Nationwide	No	−0.02
Canada		
Toronto	Yes	0.29
Nationwide	Yes	0.11
New Zealand		
Auckland	No	0.02
Waikato / Bay of Plenty	Yes	0.11
Wellington	Yes	0.55
Canterbury / Westland	Yes	0.09
Nationwide	Yes	0.11

Note: The Australian sample covers 1993:Q1–2005:Q3; the other data cover varying periods.
[a]Regression of quarterly change in the dependent variable on seasonal dummy variables.

International Evidence

Detailed tests of the impact of compositional change on median prices in other countries were not conducted due to the absence of data for the share of transactions in different segments of the market. However, in Panel C of Exhibit 4, similar tests for seasonality are analyzed for some readily available international housing price series. These are for median price series produced by the U.S. National Association of Realtors and the Real Estate Institute of New Zealand, and the mean series from the Real Estate Institute of Canada.[13] The results are comparable to Australian data, with median and mean prices in nearly all regions found to be seasonal. In some cases, seasonal dummies alone are able to explain a significant proportion of the quarterly variation in prices.

There is no compelling reason as to why pure house price changes should be seasonal. Accordingly, the results suggest that price measures in these countries are also being significantly affected by compositional change. For example, it is frequently suggested that median house prices rise in the summer in the U.S., as more large family homes sell during this period as families tend to move between school years.[14] Therefore the problem addressed in this study may be a fairly general one, suggesting that the solution proposed here may also have wider relevance for house price measures published in some other countries.

A Median-based Measure that Controls for Compositional Change

Background on Stratification

The problems illustrated in the previous section reflect the fact that the prices recorded in any quarter relate to only a sample and not the entire population of houses. Indeed, despite the significant number of transactions available each quarter, the results above suggest that the observed samples in any quarter are far from random.[15] Given that there is no ex ante way of ensuring a random sample of housing transactions, the issue becomes one of dealing ex post with the non-randomness of the sample.

The measure for the change in house prices that is proposed in this paper uses mix-adjustment, which in turn uses stratification to control for compositional change. Stratification involves dividing a population into groups (strata) such that observations within each group are more homogenous than observations in the entire population. Within each stratum, it then becomes more likely that an observed change in a characteristic of interest represents a true change rather than a spurious one due to compositional effects. Once strata have been defined, a measure of central tendency from each stratum is weighted together to produce an aggregate price measure. The resulting increase in precision from stratification will be dependent on how the strata are defined. Hansen, Hurwitz, and Madow

(1953) suggest that strata boundaries should be defined using information on all relevant variables that influence the characteristic being measured. Similarly, Lavallée (1988) notes that the most useful variables for stratifying data are those that are highly correlated with the variable of interest.

Stratification is a method employed in measuring house prices in a number of countries (Exhibit 5). However, the method used in this study to stratify the sample differs significantly from other applications in one important respect. Traditionally, the variable that has been used to group transactions is geography (Exhibit 5).[16] As mentioned earlier, grouping according to geography captures the notion that amenities and services are linked to a property's location. Another reason for grouping by location is a practical one; geographic variables are readily available in most databases of housing transactions (Goodman and Thibodeau, 2003).

This study is particularly concerned about removing the noise in median prices that result from the combination of compositional change and the extreme range in housing prices. This study is not interested in house prices across different regions of a city. Furthermore, purely geographical stratification is unlikely to divide houses into strata with the maximal feasible similarity in prices within

Exhibit 5 | Mix-Adjusted House Price Measures in Selected Countries

Index Provider	Variables
Australian Bureau of Statistics	Region
Hong Kong Monetary Authority	The saleable area of a dwelling.
Urban Redevelopment Authority (Singapore)	Dwelling type and region, with prices quoted in per square meter terms.
Bank of Canada / Royal Le Page	Region and dwelling type.
Deutsche Bundesbank / Bulwien AG	Region and dwelling type.
Ministerio de Formento (Spain)	Calculates the average price of a house per square meter. Distinguishes between dwellings based on location and size of municipalities.
Hometrack (U.K.)	Postcode and dwelling type.
Rightmove (U.K.)	Postcode and dwelling type.
Office of the Deputy Prime Minister (ODPM, U.K.)	Region, locations within region, dwelling type, old, or new dwelling and first or repeat-home buyer purchase. A hedonic equation is used to calculate the price for each stratum.

Notes: Sources include the Bank for International Settlements database, and various national sources.

strata. In particular, when cities are divided into very small areas they are likely to contain reasonably homogenous dwellings, but as areas become larger and based more on standard geographical classifications rather than economic determinants, the level of homogeneity within the group falls.[17] Accordingly, houses and suburbs are grouped into strata based on the variable that is most likely on an a priori basis to explain the price in any transaction, namely the long-term level of median prices for the suburb where the house is located.

Calculation of City-wide Quarterly Price Changes

Transactions are grouped together using the same method as outlined previously (that used in constructing the compositional change variable). That is, dwellings are first grouped by suburb, reflecting the importance of location in determining the price of a dwelling. The suburbs are then grouped into deciles (for houses in Sydney, Melbourne, Brisbane, Adelaide, and Perth) or quintiles (for houses in Canberra and apartments in Sydney and Melbourne) based on the median price of dwellings in each suburb over 2000–2004. This grouping serves to reduce the number of strata into a more manageable size (from 84–659 suburbs to 5–10 strata), contributing to simplicity of calculation, which is one of the key considerations for the simple mix-adjusted measure. It also gets around the problem that in some quarters the number of sales within a suburb may be very small (or even zero), hindering the estimation of a price movement for that suburb. Grouping suburbs in this way ensures that houses within each stratum are more likely to be similar in terms of price, and by definition ensures that we are controlling for changes in the mix of sales between higher- and lower-priced areas.

There would, of course, be many other ranking periods that could be used to group the suburbs into strata based on median prices. For example, the median price of suburbs in 1992 could be grouped to form strata for 1993, then the median prices for 1993 to form strata for 1994, and so on. However, in practice there is a very high degree of stability in the relative price rankings of suburbs: suburbs that tend to be relatively expensive in one period will tend to be relatively expensive 10 years later. Exhibit 6 illustrates this using data for Sydney, showing that the price relativities in the 2000–2004 sorting period also hold outside that period, and the Spearman rank correlations between median suburb prices in Sydney in 1996 and 2004 is 0.95. Similar results hold for other cities, for both houses and apartments. Hence, any reasonable alternative price-based strategy for ranking suburbs would result in very similar strata, and very similar estimates of price growth.

Once suburbs were grouped into deciles (or quintiles), a median price was calculated for each stratum for each quarter. The changes in the median prices for each stratum were then weighted together to calculate growth in city-wide prices. There are a number of different weighting schemes that could be used to combine these ten (or five) growth rates. The simplest method would be to take an unweighted average of the changes. This is equivalent to constructing a city-wide

Exhibit 6 | Median Decile Prices

Sydney houses, log scale, non-seasonally adjusted
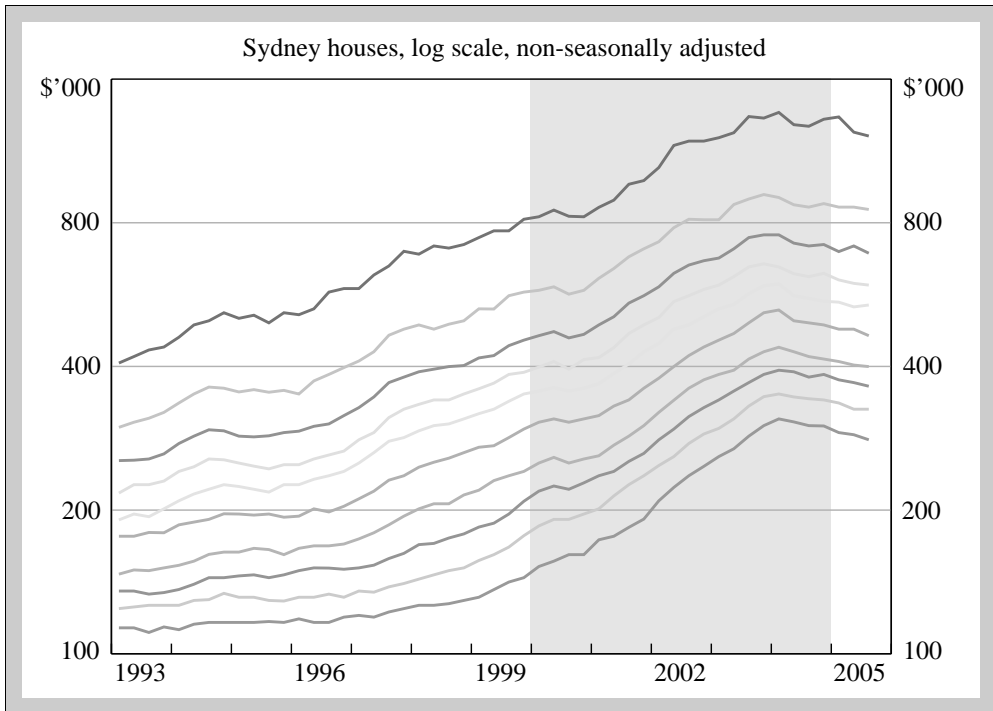
Note: The lines represent the median price for each of the 10 groups of suburbs in Sydney. The shaded area shows the period used to sort the data. The source is the authors' analysis using data from APM.

index as the unweighted geometric average of median prices in each stratum. Alternatively, the effect of using weights based on sales volumes over the 12-year period was examined, as were weights based on principal components (where changes in prices for any group can be thought as given by an unobservable city-wide movement plus an idiosyncratic component).[18]

However, different weighting schemes make very little difference to estimates of short-term price growth. In the sample, the different weightings yield measures of quarterly price changes that typically have a correlation of over 0.99. This reflects the fact that price changes in the different strata are typically quite highly correlated and that in most cases, sales volumes and principal components imply weights for each stratum that are relatively close to equal. Given that equal-weighting produces similar results to other weighting schemes and given that it is the simplest method of weighting the series, all the results shown in subsequent sections of this paper refer to the equally-weighted measure (which is labeled as the 'mix-adjusted measure').

Assessing the Mix-Adjusted Measure

The mix-adjusted measure of changes in city-wide house prices was assessed by examining how well it addressed some of the problems with conventional unstratified median measures that were highlighted previously. An additional benchmark is whether the mix-adjusted measure outperforms the change in a seasonally adjusted median price: this will indicate if the slightly greater data demands of the measure yield a significant improvement relative to accounting for seasonality in the composition of dwellings sold by just seasonally adjusting median prices. In addition, the correlation of the measure is compared with regression-based measures. Some additional perspective is then provided for the reasons for the good performance of the simple measure by comparing it to commonly used geographic stratification techniques.

Volatility

Price movements that result from compositional effects can be considered as representing noise that contributes volatility to quarterly price changes rather than being indicative of true price trends in the housing market. To assess how volatile different price series are, the quarterly movements in each of the series is compared with a measure of the 'trend' change in prices.[19] A root mean squared error (RMSE) is then calculated between quarterly growth in each of the measures and quarterly growth in the trend for each city. Since the trend measure can be thought of as capturing underlying housing price movements (or the cycle), the larger the deviations from trend, the less informative the series is about the underlying state of the housing market.

The results in Exhibit 7 indicate that in every case, simply seasonally adjusting the median price series (using the X12 program) results in a measure of price changes that is considerably less volatile than the change in the unadjusted median. However, in every case there is an additional improvement that can be gained from the simple mix-adjusted measure. Taking the reduction in the Australia-wide measure as a simple metric for the reduction in the proportion of noise in the standard median, one might conclude that seasonal adjustment can typically reduce the extent of noise by nearly 40%, but that the mix-adjusted measure results in a more significant reduction, with the average volatility falling by nearly 70%.[20]

The reduction in volatility is greatest for houses in Sydney and Melbourne. Indeed, it is noteworthy that the average deviation from trend growth in every one of the ten deciles (shown in the last column of Exhibit 7) is noticeably smaller than the average deviation from trend for the city-wide median. As a further illustration of this point, in around half of all quarters in the data sample, the quarterly change in the city-wide median lies outside the range of median price changes in all ten deciles. To be provocative, these results for Sydney and Melbourne suggest that

Exhibit 7 | Volatility in Measures of Changes in Housing Prices

	Median (nsa)	Median (sa)	Mix-Adjusted Measure	Range of RMSEs for Deciles / Quintiles
Sydney houses	4.04	2.80	1.08	1.41–2.99
Melbourne houses	4.40	2.54	1.40	1.36–3.48
Brisbane houses	1.91	1.61	1.26	1.73–4.88
Perth houses	1.90	1.49	1.07	1.73–3.10
Adelaide houses	2.20	1.37	1.27	1.71–6.07
Canberra houses	2.46	2.41	1.88	2.33–4.78
Sydney apartments	1.93	1.46	1.21	1.60–2.89
Melbourne apartments	3.45	3.01	2.11	2.78–5.59
Australian housing	2.81	1.73	0.88	

Note: The sample covers 1993:Q2–2005:Q3. The table shows deviation from trend, quarterly RMSE in percentage points.

one might get better estimates of the trend in city-wide house prices by looking at developments in a sample of only about 10% of all sales (albeit a carefully selected 10%) than from a standard city-wide median using the full sample of data.

The gains from stratification will depend on several factors including: the extent of compositional change between higher- and lower-priced properties in each city; the extent of price differences between higher- and lower-priced properties; and the extent to which the effects of compositional change can be undone via the suburb-level stratification strategy used here. The reasons for the relatively larger reduction in volatility in house price growth in the two larger cities appear to reflect both a higher degree of compositional change in these two cities (including the seasonal component shown in Panel A of Exhibit 4), and greater variation in the characteristics of the dwelling stock in Sydney and Melbourne (e.g., the median house price for the tenth decile in Sydney is on average 4.5 times higher than the median price of the first decile, compared with it being on average around 3.2 times higher for most of the other cities). In addition, since the largest cities have more suburbs, it is possible to divide these larger cities into more differentiated strata with greater variation in the average prices of suburbs in each stratum. Hence it would be expected that there would be greater gains from stratification and greater control of compositional change in the larger cities.[21]

Seasonality

By construction, the mix-adjusted measure will remove any impact on measures of price changes that results from seasonality in the composition of sales *across* strata. However, it will not control for any seasonality from compositional effects *within* strata (e.g., any possible seasonality in the composition of sales between the more and less expensive houses within any suburb). The mix-adjusted measure for the presence of any residual seasonality was tested to see if seasonality within strata is an issue. While median prices in nearly all cities were found to be seasonal, the results (available upon request) indicate that there is no evidence of identifiable seasonality in any city, nor at the nationwide level, for the mix-adjusted measure.

Comparison with Regression-based Measures

One clear advantage of a mix-adjusted measure is the simplicity of its calculation. However, more sophisticated approaches are possible, most notably hedonic and repeat-sale approaches. Of course these are not without shortcomings: hedonic regressions will only be as good as the data on housing characteristics that are available and repeat-sales estimates are likely to have significant problems in real time and can be subject to non-trivial revisions, given that estimates of price growth in any quarter will be affected by sales that occur in subsequent quarters.

Using Sydney as an example, Exhibit 8 shows quarterly growth in the mix-adjusted measure, along with estimates of growth from hedonic and repeat sales models. Growth in the seasonally adjusted and unadjusted median is also shown for comparison. The hedonic and repeat-sales estimates are taken from Hansen (2006), who uses a dataset virtually identical to the one utilized in this study. In addition, quarterly growth in the house price series produced by the national statistical agency, the Australian Bureau of Statistics (ABS), is also shown. The ABS measure is based on purely geographical stratification, using a dataset broadly similar to the one used in the current study.[22] As Exhibit 8 shows, each of the regression-based measures produces quarterly growth rates that lines up closely with those from the mix-adjusted measure. There are also noticeable differences between estimates of price growth from the city-wide median measures and the more advanced measures, indicating that it is important to adjust for compositional effects. Hansen (2006) also provides hedonic and repeat-sales estimates for house price growth in Melbourne and Brisbane. The pattern shown in Exhibit 8 also holds for Melbourne and Brisbane; the more advanced measures produce similar estimates of price growth, while growth in the city-wide median can diverge significantly from that produced by these measures.

In Exhibit 9, correlation coefficients and a measure of deviations from trend are used to compare quarterly price changes from the mix-adjusted measure to estimates from hedonic and repeat-sales price measures for Sydney, Melbourne,

Exhibit 8 | Quarterly Growth in House Prices
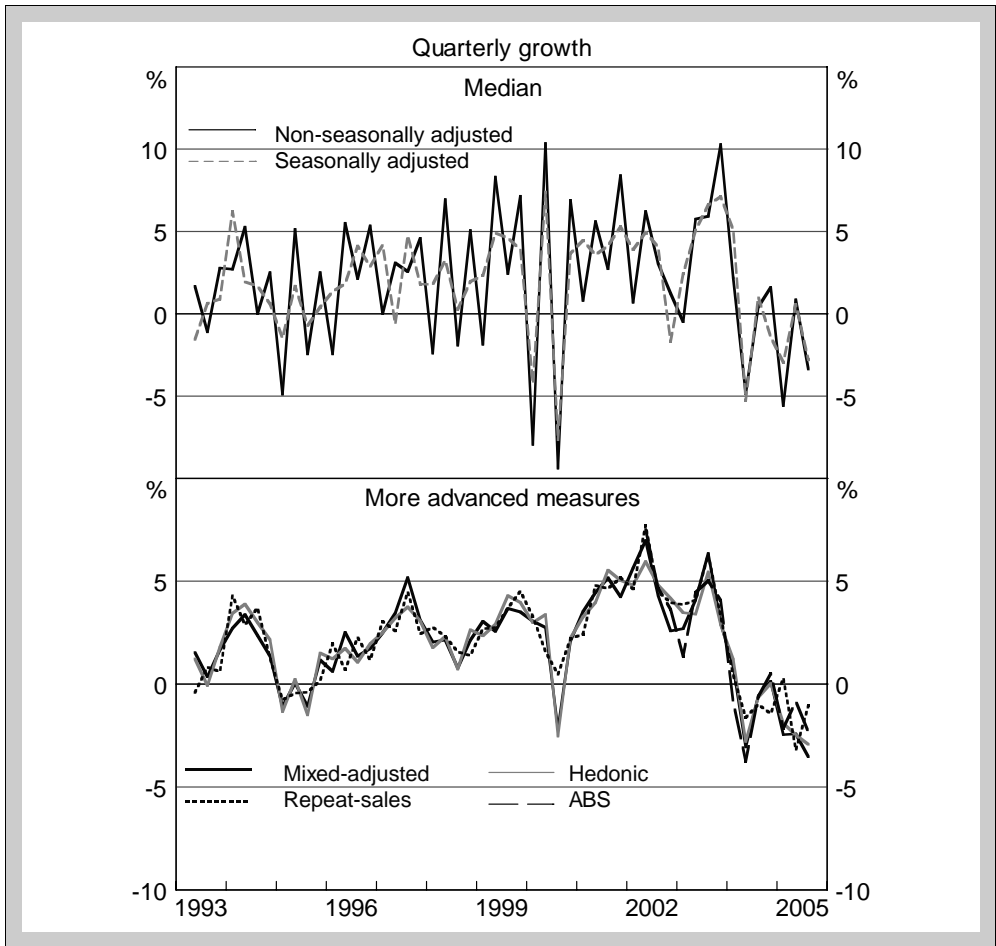
Sources: ABS; authors' calculations using data from APM.

and Brisbane. Panel A of Exhibit 9 indicates that the mix-adjusted measure of the quarterly growth in prices has a high correlation with the regression-based measures: indeed, the mix-adjusted measure tends to have a slightly higher correlation with each of the regression-based measures than the correlation between those more advanced measures. By contrast, the change in the simple median often has a fairly modest correlation with the regression-based measures, and seasonal adjustment of the median does not result in any marked increase in the correlations with the regression-based measures.

Panel B of Exhibit 9 indicates the extent to which each of the measures of house price growth deviate from a proxy of underlying house price movements. This proxy is obtained by constructing a measure of trend growth for each of the mix-

Exhibit 9 | Correlation between Various House Price Measures

	Median (nsa)	Median (sa)	Mix-Adjusted	Hedonic	Repeat-Sales
Panel A: Correlation coefficients, quarterly changes					
Sydney					
Median (nsa)	1.00				
Median (sa)	0.77	1.00			
Mix-adjusted median	0.52	0.65	1.00		
Hedonic	0.58	0.65	0.97	1.00	
Repeat-sales	0.38	0.57	0.90	0.89	1.00
Melbourne					
Median (nsa)	1.00				
Median (sa)	0.69	1.00			
Mix-adjusted median	0.65	0.71	1.00		
Hedonic	0.66	0.70	0.92	1.00	
Repeat-sales	0.42	0.57	0.76	0.69	1.00
Brisbane					
Median (nsa)	1.00				
Median (sa)	0.95	1.00			
Mix-adjusted median	0.87	0.87	1.00		
Hedonic	0.89	0.90	0.96	1.00	
Repeat-sales	0.77	0.81	0.93	0.93	1.00
Panel B: Deviation from trend (quarterly RMSE in percentage points)					
Sydney	4.11	2.95	0.97	1.02	0.86
Melbourne	4.48	2.64	1.40	1.25	1.57
Brisbane	1.96	1.69	1.25	1.25	1.03

Notes: Correlation coefficients and RMSEs across the various measures of quarterly price growth were calculated over 1993:Q2–2005:Q3. The data vintage used to calculate the hedonic and repeat-sales measures in Hansen (2006) does not correspond precisely with that used to calculate the mix-adjusted median here. In addition, Hansen uses data from a different source (the Real Estate Institute of Victoria) to calculate the repeat-sales measure for Melbourne, so the results across measures are not fully comparable for Melbourne.

adjusted, hedonic, and repeat-sales measures (using the moving-average approach outlined in Endnote 8) and then averaging these three trends. Confirming the earlier results in Exhibit 7, changes in the median and seasonally adjusted median are quite volatile with relatively high RMSEs. In contrast, the mix-adjusted measure and the two regression-based measures provide estimates of underlying house price movements that are comparable in terms of their apparent noise.

It is reassuring that the estimates of the change in house prices derived from the mix-adjusted measure are similar to those from regression-based measures,

suggesting that simple, but targeted, stratification techniques can control for a significant proportion of compositional change. However, this result may not be especially surprising as the measure is conceptually related to a hedonic approach using location as an explanatory variable. The results in Hansen (2006) indicate that the vast majority of the explanatory power in standard hedonic regressions comes from the location of properties, which (in combination with information on average suburb-level price levels) is the variable used for stratification in the methodology in the current study.

A Comparison with Alternative Price and Geographic Stratification Strategies

The preceding analysis indicates that the mix-adjusted approach overcomes many of the problems associated with unstratified median measures. A major reason for the substantial improvement appears to be the particular method used to stratify transactions. By stratifying properties on the basis of the median price for their suburb, much of the compositional change in sales movements between higher- and lower-priced properties is controlled. However, other stratification strategies are possible, an obvious alternative being on a broad geographical basis.

Unit record data for Sydney is used to construct two alternative mix-adjusted measures of price changes. Two standard geographical classifications for Sydney divide the city into 49 statistical local areas (SLAs) and 14 statistical subdivisions (SSDs). Measures are constructed using both of these geographic groupings. To produce a city-wide measure of price growth, the median house price in each geographic region is weighted by the region's share of sales over the whole sample period. In order to evaluate the relative performance of the geography-based measures of price growth, the deviation (RMSE) of each measure is calculated from the trend growth series used in Panel B of Exhibit 8.

For greater comparability, some alternative price-based mix-adjusted measures are also constructed. Instead of dividing Sydney into 10 price-based groups, the suburbs of Sydney are divided into 14 and 49 groups (the same number of groups as the geographic measures) based on the ranked median price of suburbs over 2000–2004. However, to shed further light on the stratification issue, some additional price-based measures are implemented. In particular, instead of forming measures based just on 10, 14, and 49 strata, the robustness of price-based measures is assessed by using everything from 1 stratum (equivalent to the simple city-wide median) all the way up to 60 strata (each with just 10 or 11 suburbs).

The results are shown in Exhibit 10.[23] A first point to note is that price changes estimated from the geographic-based stratifications are less noisy than the simple unstratified city-wide median. The RMSEs based on the 14 and 49 groups are 1.95% and 1.62%, respectively, versus 4.70% for the city-wide median. However, the corresponding price-based stratification measures provide a significant additional improvement over the geography-based measures, with RMSEs of

Exhibit 10 | Geographic and Price-based Groupings

Source: Authors' analysis using data from APM.

1.15% and 1.14%, respectively, for the measures based on 14 and 49 groups. This provides evidence in support of grouping data on the basis of median suburb prices rather than on a geographic basis, as the former provides a better control for changes in the mix of sales between more and less expensive properties.

An important additional result in Exhibit 10 concerns the 'granularity' of stratification in the price-based measures. The line on the graph shows how the deviation from trend (RMSE) varies according to the number of strata used to calculate price growth. Simply dividing all transactions into two groups of about 330 suburbs produces notable gains over the median measure. There are further significant gains from splitting the sample into four groups, but thereafter the RMSE is fairly constant. This implies that one can get fairly comparable estimates of movements in city-wide house prices by dividing Sydney's 659 suburbs into anything from 4 to 60 groups. This result is also confirmed by correlation analysis. Therefore, the results for Sydney shown earlier are not particularly sensitive to the decision to divide suburbs into 10 groups: indeed, there is a wide range of price-based stratification schemes that yield robust results.

Therefore, it can be conjectured that the results for the other cities are also not especially dependent upon the decision to group suburbs into deciles (or quintiles). The objective was not to fit the suburbs in each city into an 'optimal' number of groups: the choice of deciles was fairly arbitrary, although in cases of smaller sample sizes (houses in Canberra, and apartments in Sydney and Melbourne), a decision was made to instead work with quintiles to avoid small sample sizes in particular strata, especially in the incomplete real-time samples. For other applications, there will no doubt be benefits to empirically testing the optimal degree of stratification, and smaller sample sizes will presumably warrant a different number of groups, but the preliminary results here suggest that a range of strategies can yield significant benefits over simple medians.[24]

Conclusion

The conventional wisdom among most housing researchers is that median price measures provide poor measures of the pure price change in the aggregate housing stock, and that regression-based methodologies provide superior measures.[25] Yet median price measures are widely cited in the press and are used in a range of countries by industry bodies, housing lenders, and sometimes government agencies. The reason is presumably that the more advanced techniques require detailed data, are typically subject to revision as data for future periods become available, are less transparent, and require the use of statistical techniques that are not widely used outside of academic circles.

This study used a dataset of around 3.5 million transactions in Australia's six largest cities to demonstrate that compositional change between higher- and lower-priced parts of cities is a major factor behind the noise in short-term movements in city-wide median prices. However, the findings show that medians can be considerably more reliable if taken from a stratified sample. Given the importance of location in determining prices, small geographic regions are grouped according to the long-term average price of dwellings in those regions, rather than simply clustering smaller geographic regions into larger geographic regions. Therefore, the method of stratification directly controls for what appears to be the most important form of compositional change: changes in the proportion of houses sold in higher- and lower-priced regions in any period. Stratifying sales in this manner produces a mix-adjusted measure of price growth that substantially improves upon standard unstratified median measures. In particular, the mix-adjusted measure of price growth is considerably less volatile and not subject to seasonality. In addition, it is highly correlated with estimates of price growth from measures based on regression-based approaches.

The choice between different methodologies for estimating changes in aggregate prices will always involve a trade-off between various competing concerns including the ease of construction, the extensiveness of data requirements, and the extent to which the methodology controls for quality (see also Bourassa, Hoesli, and Sun, 2006). Given the recent sharp run-up in house prices in many countries,

one particular concern will be timeliness, especially for macroeconomic policymakers. However, one problem in this regard is that the groups with the best access to timely data are often industry or government bodies, which may be unable or unwilling to use the more sophisticated methodologies. This highlights two particular advantages with the approach proposed in this paper. First, the method used to control for compositional change is computationally simple, requiring nothing more than the use of a 'sort' command. Furthermore, the approach does not require the use of databases that contain a large amount of information on housing characteristics. It uses variables—transactions price and location—that are readily available from most housing transactions databases. The approach has now been used successfully in Australia for more than a year and a half, with good estimates of quarterly price changes feasible even with incomplete and biased initial samples.[26] Accordingly, the methodology outlined in this paper may be applicable for measuring house price growth in other countries as well.

Endnotes

[1] There is, for example, a significant literature on the question of whether central banks should 'lean against' perceived bubbles in asset markets by tightening policy, though there is reasonably general consensus that the conditions for such action will not generally be met (e.g., Gruen, Plumb, and Stone, 2005).

[2] For example, Bourassa, Hoesli, and Sun (2006, p. 81) put it this way: "Despite the fact that median house price indexes are widely available in several countries, they are prone to severe biases due to the heterogeneity of properties. Stated differently, such methods are unable to distinguish between movements in prices and changes in the composition of dwellings sold from one period to the next."

[3] For more information on repeat sales, see Case and Shiller (1987). An early reference to the theory underlying hedonic models is Rosen (1974). There are also numerous articles comparing the two methodologies, for example Haurin and Hendershott (1991), Crone and Voith (1992), and Meese and Wallace (1997).

[4] For example, see Clapp, Giaccotto, and Tirtiroglu (1991) and Clapham, Englund, Quigley, and Redfearn (2006) for evidence on the significant revisions to repeat sales estimates and the sometimes poor performance in estimating short-term price changes.

[5] For example, the U.S. National Association of Realtors release their monthly price series around 3½ weeks after the end of the month, whereas the OFHEO quarterly repeat sales index is released about nine weeks after the end of the relevant quarter (although there are instances in some other countries where indices using more sophisticated methodologies are released prior to median price series).

[6] Other papers that construct housing price indexes for Australia include Bourassa and Hendershott (1995), Rossini (2000), ABS (2005), and Hansen (2006).

[7] Through the rest of the paper, all calculations involving changes in prices use the change in the log of the price series. In cases where these are shown in a table or graph, they are the log change multiplied by 100 so as to correspond approximately to percentage changes.

[8] The trend is calculated as the change in the five-quarter-centred moving average of (the log of) the median price series. The weights in the moving average are 0.125, 0.25, 0.25, 0.25, and 0.125, which should remove any seasonality from the trend.

[9] There appears to be a downward trend in this ratio over the sample, perhaps because the growth in the city has been in suburbs relatively far from the center, which tend to be less expensive suburbs.

[10] The compositional change variable has also been regressed on the three measures of pure price changes outlined in the second half of this paper. The results indicate that there is no tendency for compositional change to be related to 'true' changes in house prices, so the results in Exhibit 3 reflect spurious compositional effects on median prices.

[11] The U.S. Census Bureau's X12 seasonal adjustment program decomposes a series into its trend-cycle, seasonal, and irregular components and provides tests for stable seasonality (whether or not seasonal factors have an effect on the series) and moving seasonality (whether or not the seasonal pattern changes from year to year).

[12] It is beyond the scope of this paper to consider the reasons for the seasonality in the composition of sales. However, in those cities where sales volumes are found to be seasonal, the seasonality comes more from variation in the sales volumes in higher-priced suburbs than in lower-priced suburbs. This would be consistent with some cities having particular 'selling seasons,' especially in higher-priced suburbs.

[13] The series for the U.S., Canada, and New Zealand refer to existing one-family homes, dwellings, and existing dwellings, respectively. The analysis uses data for 1975:Q2–2005:Q4 for the U.S., 1980:Q1–2005:Q4 for Canada, and 1992:Q1–2005:Q4 for New Zealand. The U.S. data have been converted to a quarterly series by averaging the monthly series.

[14] For example, the National Association of Realtors report on existing home sales in September 2004 notes that: "Most families with children, who typically buy more expensive homes, time their purchase based on school year considerations." In addition, the Real Estate Institute of New Zealand's Residential Market News report for March 2005 noted that: "A switch to higher end market residential property sales saw the residential property median price rise."

[15] This is consistent with work by Jud and Seaks (1994) using U.S. data that shows that significant differences in estimated price growth can result when Heckman's two-stage estimation procedure is used to account for potential problems in sample selectivity in the transactions that are observed.

[16] In addition to location, most measures that use stratification also group transactions according to dwelling type. As well as the measures in the table, a number of countries in continental Europe (including Austria, Finland, Hungary, and Portugal) make a rudimentary adjustment for quality by measuring prices in per square meter terms. Beyond this, most measures do not control for quality. This is probably because many datasets do not contain comprehensive information on dwelling characteristics.

[17] Further, natural boundaries will sometimes hinder the formation of larger homogenous groups. For example, expensive houses in a city may be clustered around a harbor or on either side of a river, but such natural barriers will often be used to divide the city into different geographic regions.

[18] Alternatively, if the intention is to measure changes in the value of the housing stock, suburb-level dwelling stock weights will be most appropriate.

[19] The trend is calculated using the moving-average approach described in Endnote 8. Two measures of trend are constructed first: one from an index version of the mix-adjusted

measure and the other using the seasonally adjusted median. The measure of trend used in the comparison in Exhibit 7 is the average of the growth rates of the two smoothed measures so as to ensure a fair 'horse race' between the measure and the seasonally adjusted measure (though the results are not sensitive to the precise calculation of trend growth).

[20] An alternative comparison is shown in the working paper version of this paper: for the nationwide price measure, the standard deviation of the mix-adjusted measure is 42% below that of the unadjusted median measure. Stephens, Li, Lekkas, Abraham, Calhoun, and Kimner (1994) do a related comparison of the reduction in standard deviation from the U.S. OFHEO repeat sales measure relative to the National Association of Realtors median measure: across their 11 cities, the median reduction is also 42%.

[21] Using unit record data for Sydney, the mix-adjusted measure is found to perform better in real time when compared with an unadjusted median. The real time data problem in Australia results from the existence of a lag between when a sale is agreed to and when this sale becomes recorded in a housing transactions database. Therefore initial estimates of price growth in any quarter are based on only a small sample of all transactions that will eventually be available for that quarter. Comparing an 'initial' estimate of price growth made with the data available one month after the end of a quarter, with that made from the latest vintage of data, the growth in the standard city-wide median is found to be revised by around 7.5 percentage points on average, compared with a considerably smaller 1.5 percentage point revision for the mix-adjusted measure.

[22] The ABS measure is only shown from 2002 when there is a series break and methodological improvements to the measure (see ABS 2005).

[23] Due to some constraints in the unit record data, the results here differ somewhat from the results in Exhibits 7 and 9. The measures in this section are constructed using unit record data that are of a different vintage and cover a different time span (1996:Q1–2005:Q3) to the data used in the rest of the paper.

[24] See Hansen, Hurwitz, and Madow (1953) and Everitt (1980) for more information on the theoretical issues in the optimal grouping of data.

[25] Some researchers (e.g., Gatzlaff and Ling, 1994; and Meese and Wallace 1997) have noted, however, that median price measures provide reasonably satisfactory measures of long-term price trends, and others (e.g., Crone and Voith, 1993) have sometimes found big differences between results from more sophisticated measures.

[26] The methodology is now used by Australian Property Monitors (www.apm.com.au).

References

Australian Bureau of Statistics (ABS). Renovating the Established House Price Index. ABS Information Paper No. 6417.0, 2005.

Berg, L. Price Indexes for Multi-Dwelling Properties in Sweden. *Journal of Real Estate Research*, 2005, 27:1, 47–81.

Bourassa, S.C. and P.H. Hendershott. Australian Capital City Real House Prices, 1979–1993. *Australian Economic Review*, 1995, 28:3, 16–26.

Bourassa S.C., E. Cantoni, and M. Hoesli. Spatial Dependence, Housing Submarkets, and House Prices. FAME Research Paper Series No. 151, 2005.

Bourassa, S.C., M. Hoesli, and J. Sun. A Simple Alternative House Price Index Method. *Journal of Housing Economics*, 2006, 15:1, 80–97.

Can, A. GIS and Spatial Analysis of Housing and Mortgage Markets. *Journal of Housing Research*, 1998, 9:1, 61–86.

Can, A. and I.F. Megbolugbe. Spatial Dependence and House Price Index Construction. *Journal of Real Estate Finance and Economics*, 1997, 14:1, 203–22.

Case, K.E. and R.J. Shiller. Prices of Single-Family Homes since 1970: New Indexes for Four Cities. *New England Economic Review*, 1987, September/October, 45–56.

Clapham, E., P. Englund, J.M. Quigley, and C.L. Redfearn. Revisiting the Past and Settling the Score: Index Revision for House Price Derivatives. *Real Estate Economics*, 2006, 34: 2, 275–302.

Clapp, J.M., C. Giaccotto, and D. Tirtiroglu. Housing Price Indices Based on all Transactions Compared to Repeat Subsamples. *Journal of the American Real Estate and Urban Economics Association*, 1991, 19:3, 270–85.

Crone, T.M. and R.P. Voith. Estimating House Price Appreciation: A Comparison of Methods. *Journal of Housing Economics*, 1993, 2:4, 324–38.

Everitt, B. *Cluster Analysis*. Second edition. New York: Social Science Research Council by Heinemann Educational Books, Halstead Press, 1980.

Gatzlaff, D.H. and D.C. Ling. Measuring Changes in Local House Prices: An Empirical Investigation of Alternative Methodologies. *Journal of Urban Economics*, 1994, 35:2, 221–44.

Goodman, A.C. and T.G. Thibodeau. Housing Market Segmentation and Hedonic Prediction Accuracy. *Journal of Housing Economics*, 2003, 12:3, 181–201.

Gruen, D., M. Plumb, and A. Stone. How Should Monetary Policy Respond to Asset-Price Bubbles. *International Journal of Central Banking*, 2005, 1:3, 1–31.

Hansen, J. Australian House Prices: A Comparison of Hedonic and Repeat-Sales Measures. Reserve Bank of Australia Research Discussion Paper No 2006-03, 2006.

Hansen, M.H., W.N. Hurwitz, and W.G. Madow. *Sample Survey Methods and Theory: Volume I*, New York: Wiley, 1953.

Haurin, D.R. and P.D. Hendershott. House Price Indexes: Issues and Results. *Journal of the American Real Estate and Urban Economics Association*, 1991, 19:3, 259–69.

Jud, G.D. and T.G. Seaks. Sample Selection Bias in Estimating Housing Sales Prices. *Journal of Real Estate Research*, 1994, 9:3, 289–98.

Lavallée, P. Two-Way Optimal Stratification Using Dynamic Programming. Proceedings of the Survey Methods Section. Eleventh edition. American Statistics Association, Alexandria, 1988, 646–51.

Meese, R.A. and N.E. Wallace. The Construction of Residential Housing Price Indices: A Comparison of Repeat-Sales, Hedonic-Regression, and Hybrid Approaches. *Journal of Real Estate Finance and Economics*, 1997, 14:1–2, 51–73.

Rosen, S., Hedonic Price and Implicit Markets: Product Differentiation in Pure Competition, *Journal of Political Economy*, 1974, 82(1), 34–55.

Rossini, P., Estimating the Seasonal Effects of Residential Property Markets—A Case Study of Adelaide. Paper presented at Sixth Annual Pacific-Rim Real Estate Society Conference, Sydney, January 24–27, 2000.

Stephens, W., Y. Li, V. Lekkas, J. Abraham, C. Calhoun, and T. Kimner. Conventional Mortgage Home Price Index. *Journal of Housing Research*, 1995, 6:3, 389–418.

Straszheim, M.R. *An Econometric Analysis of the Urban Housing Market*. New York: National Bureau of Economic Research, 1975.

Nalini Prasad, Reserve Bank of Australia, NSW, Australia 2001 or prasadn@ rba.gov.au.

Anthony Richards, Reserve Bank of Australia, NSW, Australia 2001 or richardsa@ rba.gov.au.