Using Hedonic Regression to Assess the Housing Rentals Component of the New Zealand Consumers Price Index

Paper presented at New Zealand Association of Economists 50th Anniversary conference,
at Wellington, New Zealand,
1 – 3 July 2009

Frances Krsinich
Senior Research Statistician, Prices Unit, Statistics New Zealand
P O Box 2922
Wellington, New Zealand
frances.krsinich@stats.govt.nz
www.stats.govt.nz

Abstract

The housing rentals index is the highest weighted component of the Consumers Price Index, contributing nearly 8 percent as at the June 2008 quarter. The estimation of non-Housing New Zealand rents is based on data from an area-based sample survey, updated with new bonds lodged with the Department of Building and Housing data. A 'matched sample' approach is used to remove the effect of changing sample composition, and therefore changing quality characteristics, from the measurement of price change. Although matched sample approaches for price change measurement are common across official statistics agencies, there is a potential risk that some pure price change implied by the changing population is removed.

As part of a more general review of the housing rentals index estimation, we used hedonic regression models to examine whether there is any residual pure price change being smoothed out of the index due to the matched sample approach. In the process, we explored different approaches to the specification of the hedonic regression models and the corresponding index calculation from the model parameters.

This paper will present the results, and give initial conclusions.

Introduction

We present the results of our investigation of two questions:

1. Is the restriction of the sample to those that exist in the current and previous quarter (ie the ‘matched sample’) biasing our estimation of the quarterly movement in rentals?

2. How effective, generally, is the current estimation method at controlling for the effect of compositional change in the sample over time?

We show that:

- The restriction of the sample is not biasing the price measurement to a level of any practical significance.

- The current estimation method does well at controlling for compositional change.

The rental data is longitudinal – that is, we have repeated rental information for dwellings in the sample, and this enables us to control for unobserved characteristics in a way not usually possible in hedonic estimation. This has given us the opportunity of comparing hedonic indexes controlling for observed versus all characteristics fixed at the dwelling level.

The results to date suggest that the current estimation method is doing a good job of controlling for compositional change in the measurement of rental movements. The next major step in ensuring the good quality of the rental estimation will be to consider whether the sample needs to be recalibrated to population benchmarks and/or reselected.

Background

'Actual rentals' contribute almost 8 percent of expenditure weight of the Consumers Price Index, as at the 2006 rebase.

Table 1

Contribution of 'Other rentals' to the Consumers Price Index

Level	Description	Percent
All groups	Consumers Price Index	100
Group	Housing and household utilities	22.75
Subgroup	Actual rentals for housing	7.85
Class	Actual rentals for housing	7.85

Note: 'Other rentals' contribute approximately 90 percent of the 'Actual rentals for housing' class, which also contains Education accommodation and Housing New Zealand rentals.

Table 2

Expenditure Weights for Actual Rentals

Rebase	1999	2002	2006	2008
Expenditure weight	6.16%	5.54%	6.87%	7.85%

Sample design and estimation

We construct the rental index from the movement in average rent for dwellings with different numbers of bedrooms (1 bedroom, 2, 3 and 4 or more bedrooms) and regions (Auckland, Wellington, Other North Island, Christchurch, and Other South Island). The dwelling size and region combinations are weighted together with expenditure weights that are fixed between CPI reweights (every three years).

The initial sample of the current design (started in 1998) was a stratified random sample of areas (Primary Sampling Units (PSUs)). The PSUs were stratified by low, medium and high average rent based on Census of Population and Dwellings data. A scoping questionnaire was sent to all dwellings to determine if they were rented, and then each rental dwelling in the sample was surveyed every quarter.

Currently the sample size (of rented dwellings) is approximately 2,200.

Given possible changes in ownership of dwellings, from rented dwelling to owner-occupied or vice-versa, new bond lodgements with the Department of Building and Housing are used to identify when dwellings within the sampled meshblocks become in scope of the survey (either when an established dwelling starts being rented or when a newly constructed rented dwelling comes onto the market). When such dwellings are identified, they are enrolled into the survey. While the process of identifying new rental dwellings within surveyed meshblocks was introduced in 1999, administrative changes caused the process to lapse in late 2001 and it was not reinstated until the June 2006 quarter. The number of births each quarter varies between about 120 and 220.

Each quarter, there is approximately 10 percent non-response, and rents are imputed by carrying forward the last known rental value. For persistent non-response, the value is carried forward for five quarters before the non-respondent is assumed to be a 'death' and is removed from the sample.

The estimation method was modified in 2000, when a 'matched sample' approach was introduced to ensure that differences in the composition of the sample due to new rental dwellings entering and old rental dwellings exiting would not contaminate the estimation of price movement. So, each quarter, the movement is calculated from the dwellings that exist in both the previous and current quarters only.

Figure 1 shows the number of matched, births, deaths and otherwise unmatched in the sample for each pair of adjacent quarters since quarter three of 2006, which is when the process for introducing births to the sample had stabilised. Dwellings existing in both quarters can also be 'unmatched' if the number of bedrooms change – as shown in figure 1 there is a negligible number of these.

Figure 1

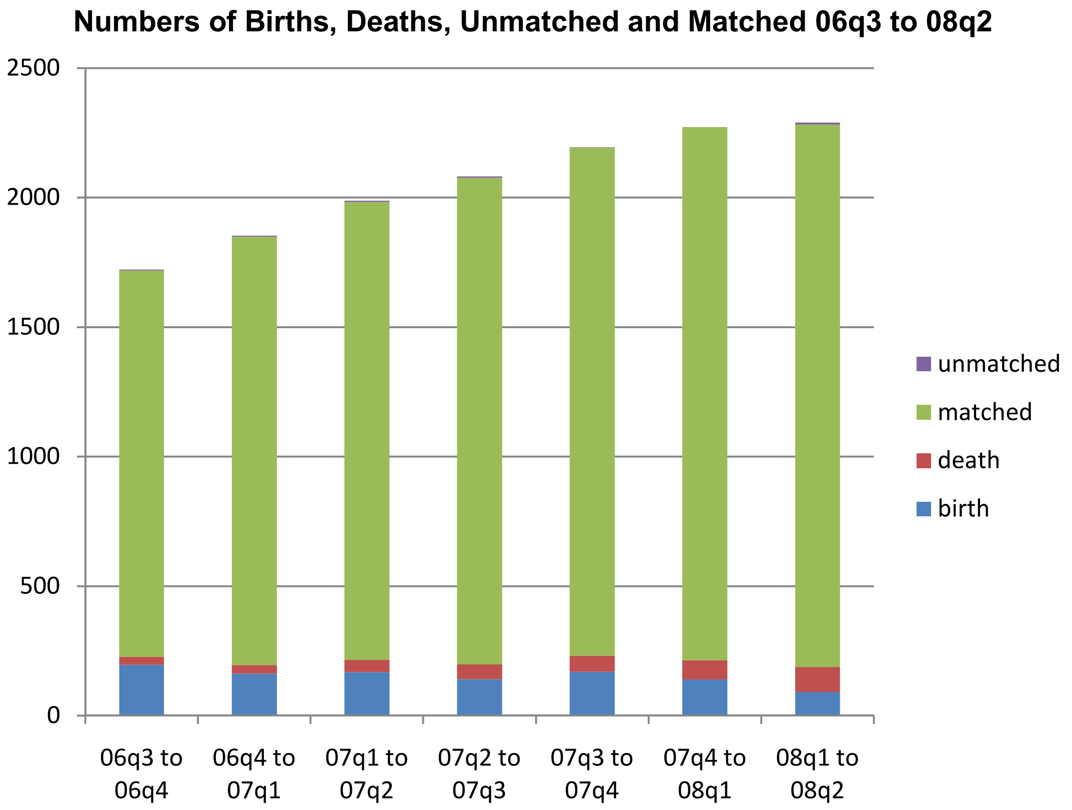

Matched sample estimation

In each quarter, average rents are calculated within strata defined by number of bedrooms and CPI area (ie five broad regions), based on those rental dwellings in the matched sample – ie those that exist in the previous and current quarter, with unchanged characteristics.

Consider the following simple example of dwellings within a given strata – say, three bedroom rental dwellings in Christchurch:

Table 3

Matched Sample Estimation of Price Change

Dwelling id	Rent_Q1	Rent_Q2	Rent_Q3
1	350	350	
2	270	270	270
3	150	160	160
4	280	280	285
5	120		
6			280
7		310	310

In this example, dwelling 1 and 5 drop out of the sample (or 'die') in quarter three and two respectively, while dwelling 6 and 7 are new to the sample (ie they are 'born') in quarters three and two respectively.

In quarter two, we would calculate the movement from quarter one to quarter two from those in the matched sample, that is, dwellings 1, 2, 3 and 4. So, the movement would be (350+270+160+280)/4 divided by (350+270+150+280)/4 = 265/262.5= 1.0095 – a 0.95 percent increase in average rent.

Note that sampling weights would be incorporated in the actual calculation. For simplicity, we omit weights from this example.

Similarly, the estimate of the rental movement for quarter two to quarter three would be based on dwellings 2,3,4 and 7 only, so the estimated rent movement would be 256.25/255 1.0049 – a 0.49 percent increase in rent.

If the averages were calculated from the full sample in each quarter, the estimates would have been an increase of 17.09 percent and a decrease of 4.74 percent, respectively.

Obviously this is an extreme example. In reality, the proportion of unmatched dwellings in each pair of quarters is relatively small, but it illustrates the motivation behind the matched sample approach, which is to compare like with like when calculating price movements.

Figure 2 shows the difference in the 'Other rentals' index calculated from the full data in each pair of quarters – *current method (full)*, versus the matched sample – *current method (matched)*. The discontinuity between 06q2 and 06q3 corresponds to the introduction of new procedures for introducing births to the sample, which resulted in a backlog of births being added to the sample in quarter two of 2006.

As can be seen, the index calculated using a matched sample is both flatter and less volatile than that calculated from the full data.

Figure 2

The Current Estimation Method on Full and Matched Samples

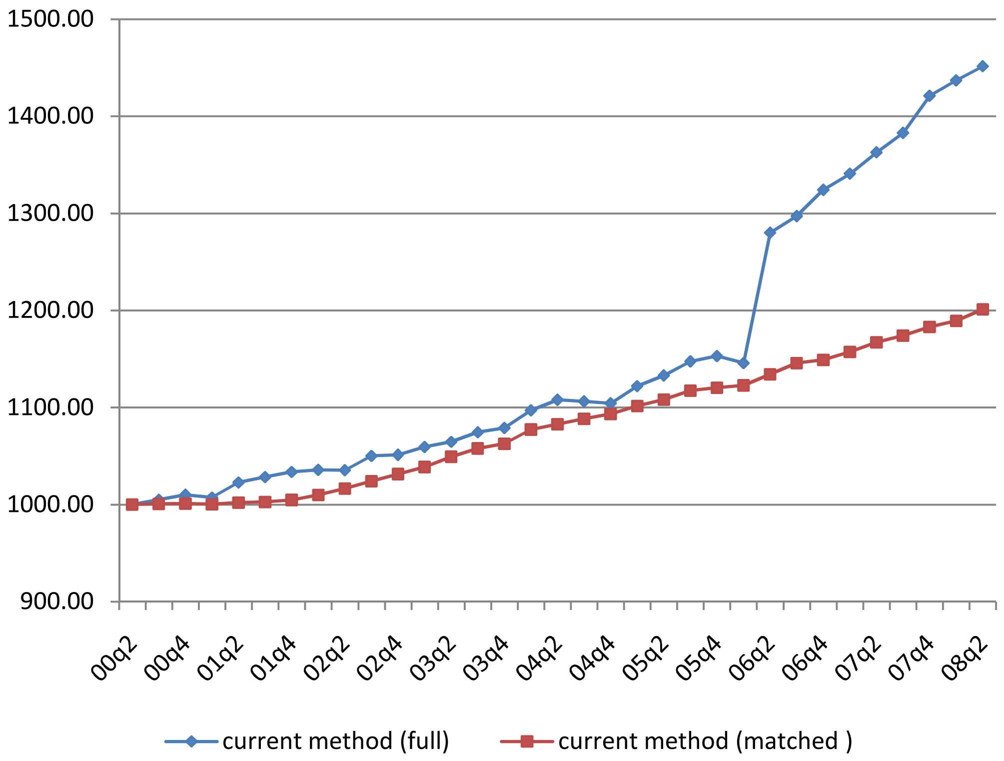

Concerns about the matched sample approach

Concerns were raised about the matched sample approach (Smith 2008).

New dwellings to the sample will not contribute to the estimation in the first quarter that they are observed in the sample, as they won't be part of the matched sample. The first rent movement contributed by new dwellings will be that from the first to second quarters that they are observed. Given that rent movements tend to be flat between tenant changes, this first observed movement is also likely to be flat. Is this then introducing a bias to the index? If the index derived from the full data – *current method (full)* in figure 2 – is a better estimate of the underlying change in rentals, then there is potentially quite a downwards bias being introduced into the rental index by our use of the matched-sample approach, which in turn would contribute downward bias into the CPI.

We already know that producing the index from the full versus the matched data can give quite different results, as shown in figure 2, but it is also clear from our example above that the price change estimated with our current estimation method will be contaminated by changes in the composition of the sample when the full data is used.

We therefore approached this investigation by creating hedonic indexes as benchmarks. These enabled us to answer two questions:

1. Is the restriction of the data to those dwellings matched in each set of adjacent quarters biasing the index?

2. How well is the current estimation method performing, generally, compared with a hedonic index?

Hedonic regression

Hedonic regression is essentially multivariate regression where price (or the log of price) is modelled as a function of time and characteristics of the entity whose price change we are estimating. Hedonic regression was introduced by Court in 1939 for the measurement of automobile prices. He said "… the basis of computation is simply *the measurement of the relation of price to time, holding usefulness (as reflected by specifications) constant*. The statistical technique involved is standard multiple regression procedure."

The hedonic model can be estimated from data pooled across all periods, or separately on adjacent periods, which then requires that the estimated movements be chained together. Pooling across all quarters assumes that the 'shadow prices' of characteristics are constant over time, which may be incorrect. On the other hand, using adjacent periods only may not provide enough data for a robust estimation. A compromise between these two extremes is to use a moving window of some number of periods.

The hedonic regression of the rent of dwelling i in period t on its characteristics set z_{tki} is given by

$$\ln p_{ti} = \beta_0 + \sum_{t=2}^{T} \beta_t D_t + \sum_{k=1}^{K} \beta_k z_{tki} + \varepsilon_{ti} \qquad (1)$$

where D_t is a dummy variable for the time periods (ie D_2 = 1 in period 2, 0 otherwise) and z_{tki} is the set of characteristics for dwelling i in period t, which could include both dummy variables (eg for region) and continuous variables (eg age of dwelling).

In our case, we have categorical variables for CPI area, stratum, PSU and number of bedrooms for each rental dwelling.

If the hedonic regression is based on each set of adjacent periods separately, the formula becomes

$$\ln p_{ti} = \beta_0 + \beta_c D_c + \sum_{k=1}^{k} \beta_k z_{tki} + \varepsilon_{ti} \qquad (2)$$

Where $D_c = 1$ in the 'current' (ie later) period of the pair of adjacent periods, or 0 in the previous period.

The hedonic index is then derived from the parameters on time. The price movement from period t to period $t+n$ is

$$\frac{P_{t+n}}{P_t} = \frac{\exp(\beta_{t+n})}{\exp(\beta_n)} \qquad (3)$$

Limitations and opportunities of the rental data

Limited observed characteristics

The rental data is quite limited in the characteristics of the rental dwellings available. We have: available number of bedrooms (1 ,2 ,3, and 4 or more); broad region (Auckland, Wellington, Other North Island, Christchurch, and Other South Island); stratum (high, medium and low rent – at the PSU level, and defined on census data); PSU indicators (PSU is the 'primary sampling unit' – an aggregation of meshblocks); and meshblock (a fine-level regional variable – on average around 100 dwellings).

Unobserved characteristics will be controlled only to the extent that they are correlated with one or more of the observed characteristics. So, for example, if the average quality of rental dwellings in a selected meshblock is increasing over time (after controlling for number of bedrooms), then a standard hedonic index, including all the characteristics we have access to, will be biased by this quality change.

Quantity weights

Unlike many hedonic indexes, which are constructed from prices from catalogues and do not have quantity weights, our weighted sample data is representative of the quantities of different rental dwellings of particular characteristics in the population of rented dwellings.

Longitudinal data

This longitudinal structure gives us quite an opportunity because, by using fixed effects regression, we can control for all characteristics that are fixed at the dwelling level. That is, we can control for both observed *and unobserved* fixed characteristics.

Fixed effects regression

There must be two or more measurements on the same dependent variable for each individual in the sample and, for at least some of the sample; the values of the independent variables of interest must vary across at least two of the measurement periods.

Each individual serves as their own control – this is done by making comparisons within individuals (in our case rental dwellings) and then averaging those differences across all the individuals in the sample.

Fixed effects methods do not allow estimation of coefficients for variables that have no within-individual variation such as, in the case of rental dwellings, CPI area, stratum, PSU or meshblock, but we are ultimately only interested in the parameters on time to calculate the rental index.

For non-experimental data such as a sample surveys, fixed effects methods tend to reduce bias at the expense of greater sampling variability. The rental sample is large enough to cope with this increase in sampling variability.

The fixed effects hedonic model can be formulated as follows:

$$\ln p_{ti} = \beta_0 + \sum_{t=2}^{T} \beta_t D_t + \sum_{k=1}^{K} \beta_k z_{tki} + \alpha_i + \varepsilon_{ti} \qquad (4)$$

where the variables have the same meaning as in formula (1), with the addition of α_i which represents all the dwelling-specific variation that is fixed over time.

An example of an unobserved characteristic that is fixed across time at the dwelling level is size (after controlling for number of bedrooms, which is observed).

However, note that this model will *not* control for unobserved characteristics which *vary over time* at the dwelling level. For example, renovations or deterioration due to aging. Although, note that these two particular unobserved characteristics will affect the estimation in opposite directions. See the ‘Further research’ section for more discussion on the issue of dwelling aging. Any *systematic* change in the composition of the sample with respect to unobserved time-varying (at the dwelling level) characteristics will introduce bias into the index.

Following Allison (2005):

Ordinary Least Squares (OLS) will produce optimal estimates of the parameters, but including dummy variables for the α_i terms is computationally tedious. We get identical results by 'conditioning out' the α_i terms and performing OLS on deviation scores. We compute the means over time for each dwelling and each time-varying variable – both response and predictors.

$$\overline{y_i} = \frac{1}{n_i}\sum_t y_{it} \quad (5)$$

$$\overline{x_i} = \frac{1}{n_i}\sum_t x_{it} \quad (6)$$

where n_i is the number of measurements for dwelling i.
In our case the response y is logged rent, and x is the number of bedrooms (all other observed characteristics are time-invariant).

Then we subtract the dwelling-specific means from the observed values of each variable:

$$y^*_{ti} = y_{ti} - \overline{y_i} \quad (7)$$

$$x^*_{ti} = x_{ti} - \overline{x_i} \quad (8)$$

Then y^* is regressed on x^* and the time dummies. Using SAS's PROC GLM with the ABSORB command achieves all this.

Results

We initially calculated the index for rentals using a 'standard' hedonic regression approach, which controlled for only the observed characteristics of dwellings in the data, and did not incorporate the inter-dwelling correlation. Estimating the quarterly movements from adjacent quarters, we compared the index calculated from the full sample to that from the sample restricted to the matched sample in each pair of adjacent quarters.

We then used fixed effects, to control for all characteristics (both observed and unobserved) fixed at the dwelling level. With this fixed effects hedonic estimation, we then tested whether the restriction of the sample to that in the 'matched' subsample is biasing the index. The nature of the fixed effects estimation is such that we could not directly compare the full to the matched sample for an index chained from adjacent-quarter estimations, so we estimated an upper bound on the potential bias introduced by the restriction to the matched sample, by using a 'matched for four quarters' definition with pooled yearly models.

Standard hedonic regression on adjacent quarters

We calculated the standard hedonic regression on adjacent quarters, which allows the parameters on characteristics to change over time, and then chained the resulting movements together to get the index for the full time period – quarter two of 2000 to quarter two of 2008.

By restricting the sample to those that are matched for each pair of adjacent quarters (ie they exist in both quarters and have the same number of bedrooms), we calculated the corresponding standard hedonic regression restricted to the matched sample.

Figure 3 shows the results compared with the index from the current estimation method on both the full and matched samples. The hedonic adjacent quarter index based on the full sample – *hedonic adj q (full)* – is flatter than the index calculated using the current method on the full sample – *current method (full)* - which implies that the current method applied to the full sample is biased by compositional change in terms of the observed variables (not to mention unobserved characteristics).

The hedonic index calculated from the full samples for each pair of adjacent quarters is quite different from that calculated from the matched samples. On the face of it, this seems to imply that the restriction to a matched sample is biasing the rental estimation downwards. However, at this point we note that we are controlling for only a limited set of characteristics (number of bedrooms and various aggregations of region: PSU, CPI area, and stratum). In addition, we are not properly incorporating the longitudinal structure of the data into the estimation, so there is correlation within dwellings that is not being taken account of in the estimation.

Figure 3

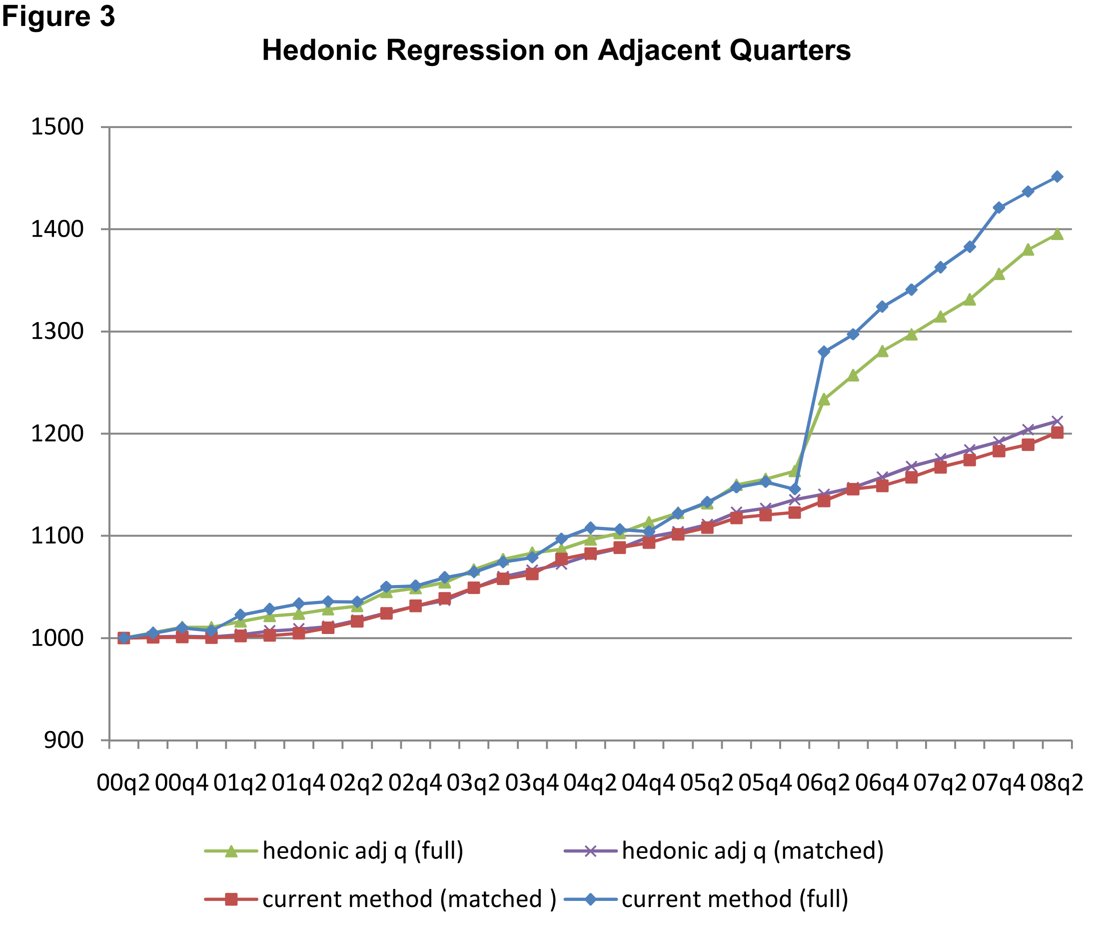

Fixed effects hedonic regression

To better estimate the rental index by using fixed effects hedonic regression, we started by using the data pooled across all time periods from quarter two of 2000 to quarter two of 2008. This holds the parameters for characteristics fixed across the observation window and will be biased in as much as the composition of the sample changes systematically over time in terms of these characteristics. However, it was a useful place to start to get an indicative fixed-effects index.

Figure 4

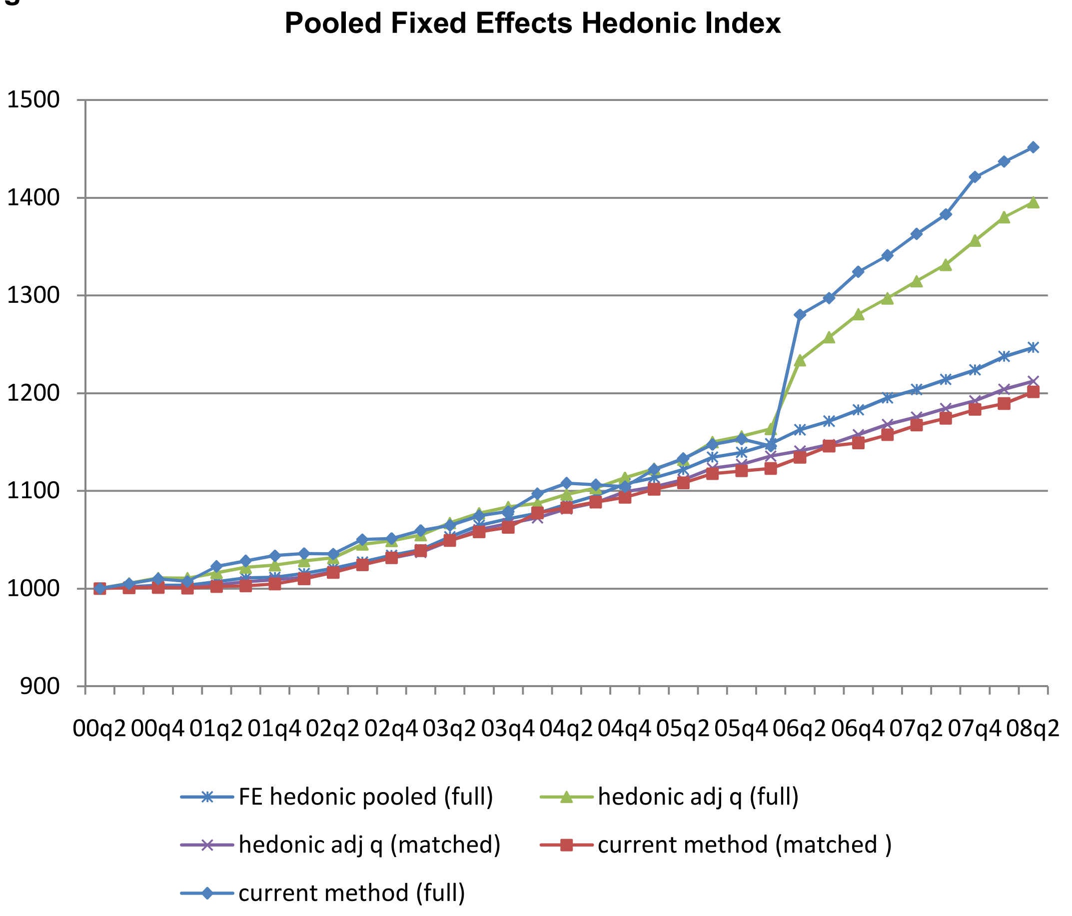

As shown in figure 4, the fixed effects hedonic index based on the full data – *FE hedonic pooled (full)* – is significantly flatter than either the standard hedonic or the current method based on the full data. This implies that there is unobserved quality change in the sample that is not controlled for by either the current method on the full sample or the standard hedonic. It is interesting to see how much closer this fixed-effects hedonic index is to the currently estimated index using the matched sample – *current method (matched)*.

The next step is to explicitly test for any biasing of the index via the restriction to the matched sample in each pair of adjacent quarters, and to allow the parameters estimated for characteristics to change over time. Note, though, that the only time-varying characteristic we observe, and therefore estimate parameters for, is number of bedrooms. There is relatively little variation in the data for this variable as most rental dwellings retain the same number of bedrooms over time.

Comparing fixed effects on the full versus matched samples

Initially, we estimated the fixed effects hedonic index on adjacent quarters, comparing the index calculated from the full sample to that from the matched sample. However, since fixed effects regression requires at least two observations for each dwelling, our 'full' data for each pair of adjacent quarters was just the matched sample plus any dwellings with a change in number of bedrooms. As noted above, there are very few dwellings whose number of bedrooms changes.

This could be seen as a demonstration of the point that Silver and Heravi (2003) make, that the adjacent quarters hedonic method (which they refer to as the 'exact hedonic approach') is in fact closely related to the matched-sample method proposed 'on pragmatic grounds' by Turvey (1999a,b).

To work around this limitation we fitted the fixed effects hedonic regression for four quarters at a time, and incorporated a more stringent definition of 'matched' to get an upper bound for the potential bias that we were interested in. So, we subset the full data to a 'yearly-matched' subset where the dwelling exists in all four quarters and has the same number of bedrooms for all four quarters. If the difference between the fixed effects hedonic index calculated from the full data versus that calculated from the 'yearly matched' subset of the data are very similar, then we will have shown that the restriction to the less stringent 'quarterly-matched' subset does not bias the index estimation.

Figure 5 shows that the restriction of the sample to the subset which is matched across the year has an almost imperceptible effect on the resulting fixed-effects hedonic index. Therefore, we are satisfied that the restriction to the matched sample of the data used for calculating the rental index does not, in itself, introduce any bias to the index.

Figure 5

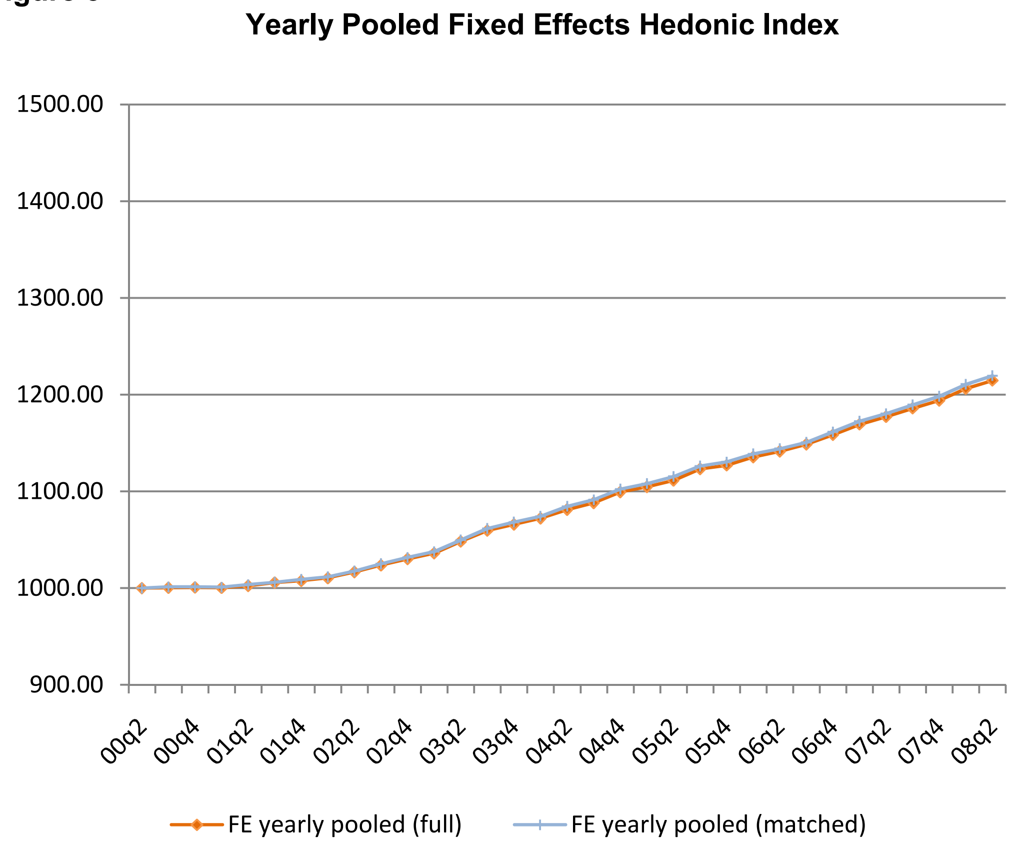

All results compared

Figure 6 shows all the indexes together for comparison.

Figure 6

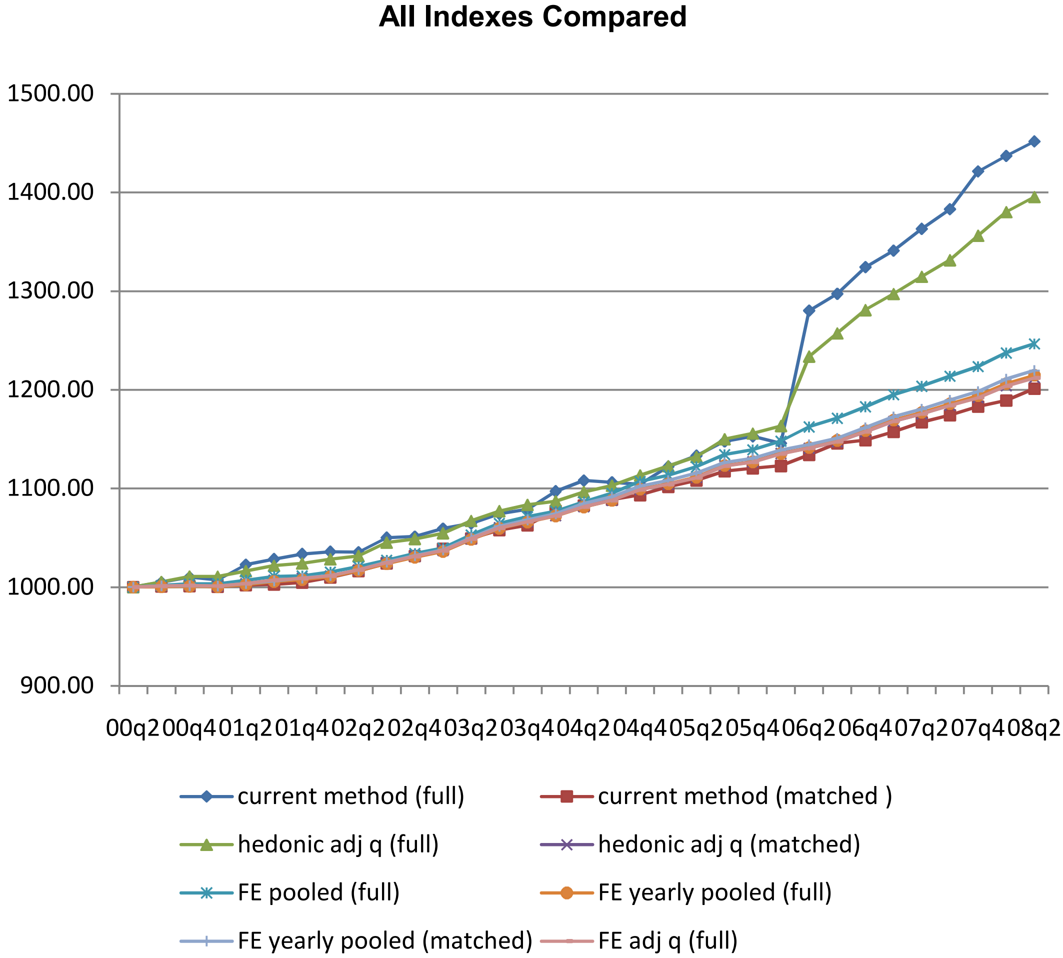

Current method (matched) is the rental index as currently estimated from a matched-sample estimation of average rents within strata. *Current method (full)* shows what the index would look like if the current method was used to estimate the index based on the full sample rather than the matched sample in each pair of adjacent quarters. These indexes are significantly different, which is what motivated the concerns about the use of a matched-sample approach.

Hedonic adj q (full) and *hedonic adj q (matched)* show the index calculated from the full and matched samples, respectively, where quarterly movements are calculated from successive pairs of adjacent quarters and chained together. These are based on 'standard' hedonic regressions, which control only for the observed characteristics in the data and do not correct for the within-dwelling correlation in rents. The difference between these is smaller than the corresponding difference between using full and matched data with the current estimation method, but is still significant, which suggests that there is bias due to the restriction to the matched sample and/or that the hedonic regression is not well specified. The discontinuity in the 06q2 to 06q3 movement of the *Hedonic adj q (full)* index (when the backlog of births was incorporated) suggests that this hedonic index does not control well for changes in the characteristics of the sample.

The fixed effects hedonic indexes *FE yearly pooled (full)*, *FE yearly pooled (matched)* and *FE adj q (full)* are so similar as to be almost indistinguishable in this graph. As explained above, the difference between *FE yearly pooled (full)* and *FE yearly pooled (matched)* gives an upper bound on the potential biasing due to the restriction of the data to a matched sample. *FE adj q (full)* is included to give an indication of the difference to the index when the parameters for characteristics are held constant in four-quarter windows (for the *FE yearly pooled(full, matched)* indexes) rather than being allowed to change with each new quarter-to-quarter estimation (*FE adj q (full)*). Clearly, this restriction has no practical effect on the index.

FE pooled (full) is the index calculated from the fixed-effects hedonic regression on the data pooled across all eight years. The difference between this index and the FE adj q (full) index is the effect of holding the parameters for characteristics constant across all eight years of data. This difference shows that we should be allowing the parameters to change over time, although yearly seems sufficient as shown above.

Conclusions

Our use of hedonic indexes as a benchmarking tool has answered both our initial questions:

- The restriction of the sample to the subset that exists in both the current and previous quarter (ie the 'matched sample') is not biasing our estimation of the quarterly change in rentals. This is shown by the similarity between *FE yearly pooled (full)* and *FE yearly pooled (matched)* in figure 6.

- The current estimation method is doing a good job of controlling for the effect of compositional change in the sample over time (in terms of characteristics fixed over time at the level of the dwelling). This is shown by the similarity between *FE adj q (full)* and *current method (matched)* in figure 6.

The next area for attention will be an assessment of the representativeness of the rental sample. This will involve, at the least, a recalibration of the sample weights to updated population benchmarks and, possibly, a full reselection of the sample. Although the sample is being continually updated with births from bond data, there is potential for respondent fatigue over time to lead to non-response bias.

While the immediate objective of this investigation was to assess the performance of the current estimation method, a number of areas for further research in hedonics have been raised. The use of hedonic regression is an area that we are interested in, generally, both as a benchmarking tool to assess the performance of non-hedonic approaches to controlling for change in quality, and in production. We have been using a hedonic method for the estimation of used car price change in the index for the past eight years (Krsinich, 2000), which we plan to re-evaluate shortly. Also, potential future utilisation of scanner data (ie retail transaction data) for price measurement is likely to require the use of hedonic methods.

Further research

Two areas for further research that have arisen during the course of this investigation are the price effect of the aging of rental dwellings, and the potential for adjusting for spatial correlation.

Aging of the sample

Each dwelling in the sample is aging over time and, by construction, we know that the current matched sample estimation would be biased to the extent of any price decrease associated with this aging. We conclude that the practical effect of this biasing must be small, because the current method gives similar results to the hedonic index as shown above. The hedonic index will be unbiased by aging at the individual level, as it is estimated from a sample that is designed to be representative of the current population of dwellings at each quarter, due to the continual birthing into the sample of new rental dwellings and removal of dwellings that are no longer rented. This similarity between the two indexes may be due, in part, to biasing of the matched sample estimation in the opposite direction through a similarly uncontrolled 'renovation' effect. More thought about the effect of aging is required.

This longitudinal rental data does give us the potential to estimate the price effect associated with the aging of rental dwellings. Although, it is likely to require some assumptions about constant proportional effect of aging over time (that is, that the proportional price effect of aging five years is the same for a 10-year-old dwelling as for a 20-year-old dwelling). There is scope for further research here.

Spatial correction

In their estimation of house price indexes for the Sydney region, Hill, Melser and Syed (2009) incorporated an adjustment for spatial correlation by using a spatial weight matrix, which identifies the neighbours of each observation based on longitude and latitude of each property.

We have gridpoint information for meshblocks. Gridpoints are a two-dimensional projection of longitude and latitude. Although this is spatial information at a broader level than that used by Hill et al, we believe there is scope for further research comparing results using spatial correction with fixed effects and/or a combination of both.

References

Allison, P (2005). Fixed effects regression methods for longitudinal data using SAS, SAS Institute, Cary, NC, USA

Court, A T (1939). "Hedonic price indexes with automotive examples" in *The Dynamics of Automobile Demand*. The General Motors Corporation, New York, 99-117.

Hill, R, Melser, D & Syed, I (2009). "Measuring a boom and bust: The Sydney housing market 2001 – 2006" (draft)

ILO/IMF/OECD/UNECE/Eurostat/The World Bank (2004). Consumer price index manual: Theory and practice, International Labour Office, Geneva, chapters 7 and 19

Krsinich, F (2000). "The hedonic regression parameter method for calculating the Used Cars index" Internal Statistics New Zealand research paper.

Silver, M & Heravi, S (2003). "The measurement of quality-adjusted price changes" in Scanner Data and Price Indexes, eds Feenstra, R & Shapiro, M, The University of Chicago Press, Chicago and London, 277 – 316.

Smith, M (2008). "Measuring rents in the CPI" Internal information paper, Reserve Bank of New Zealand.

Turvey , R (1999a). *Draft of elementary aggregate indices manual*. Available at [http://www.turvey.demon.co.uk]. Revised April.

Turvey, R (1999b). “Incorporating new models into a CPI: PCs as an example” In Proceedings of the Measurement of Inflation Conference, ed. M. Silver and D. Fenwick, 505-10. Cardiff, Wales: Cardiff Business School.