
The measurement of house prices

By Gregory Thwaites and Rob Wood of the Bank's Structural Economic Analysis Division.

House prices are an important consideration in assessing macroeconomic developments in the United Kingdom. But the special characteristics of housing—heterogeneity, infrequent sale and negotiated prices—give rise to important issues that complicate their measurement. There are several valid concepts of house prices—such as the average transaction price, the price of a typical house and the housing stock deflator—each of which is useful for a different purpose. Users must therefore be careful to match the measure they use with the concept of house prices they are interested in. Furthermore, all the available measures are volatile, so high-frequency changes in house price inflation should not be expected to persist.

Introduction

House prices rose by around 25% in 2002. This led to an increase in the value of the housing stock equivalent to around 85% of annual household disposable income, which more than matched the falls in household equity wealth. A rise in house prices increases the amount of money that new buyers and people trading up must borrow to fund house purchases, as well as the amount of equity available to existing homeowners to borrow against. The resulting mortgage equity withdrawal (MEW)—extra borrowing secured on housing that is not spent on additions to the housing stock—can be used to fund consumption, pay down other debts, or purchase financial assets. In the year to 2002 Q3, MEW amounted to £38 billion, equivalent to 5.3% of disposable income over this period. House prices also feed directly into the RPI, through measured physical housing depreciation (which represents the cost of maintaining the quality of the housing stock) and mortgage interest payments: physical housing depreciation accounted for around one third of annual RPIX inflation in December 2002. So house prices are an important consideration when assessing macroeconomic and financial developments in the United Kingdom.

But, as with many economic statistics, their measurement poses significant conceptual and practical problems. There is a range of available measures of house prices, and these can give conflicting or

misleading signals about levels of or changes in house price inflation. This article discusses how house prices are measured, and how the available measures should be interpreted.

Why are house prices difficult to measure?

Houses are both durable goods and tradable assets. They provide a flow of rents to the owner, or housing services to the occupier, and there is an active secondary market in which houses can be bought and sold. But the price of housing is harder to measure than that of most other goods and assets, because of three key distinguishing characteristics.

First, and most importantly, dwellings are heterogeneous. No two dwellings are identical, if only because they cannot occupy quite the same location. This means that we cannot always reliably predict the sales price of a given dwelling from the price of another.

Relatedly, we cannot easily observe the market price of a given dwelling without it being sold. Dwellings are typically transacted at a price reached through negotiation or at auction, so the advertised price can be a poor guide to the eventual selling price. In contrast, other goods have advertised list prices, and many financial assets are traded on exchanges with quoted bid and offer prices. Professional assessors can be employed to estimate the price of a house, but this is expensive and can be a poor guide to the eventual sales price.⁽¹⁾

(1) See Goolsby (1997) for a discussion of the systematic biases found in previous studies.

Our set of price observations is therefore usually restricted to transactions prices, so changes in the types of houses transacted will affect the mixture of prices we can observe.

Finally, houses are generally sold infrequently: over the 1990s, the number of private dwellings sold per year was around 7% of the stock. At this rate, each house would be sold on average approximately once every 14 years. So the most recent price observation for a given house will be on average 7 years old, and will therefore be an unreliable guide to the price it would fetch today.

A simple average of transactions prices in a given period has a clear interpretation: the mean price of houses sold in that period. This is a useful measure if one wants to estimate the value of turnover in the housing market (which will be related, for example, to stamp duty receipts and estate agents' turnover). But we may equally well be interested in the value of the total housing stock or the price of a representative house, in which case this simple average may be misleading.

Different types of house may be transacted at different rates or at different times. So the mean transacted price will be an unreliable guide to the mean price of all houses: changes in this mean price over time may reflect changes in the mix of houses being sold rather than in the value of the stock of dwellings. Suppose, for example, that detached houses are typically more expensive than terrace houses. If the proportion of detached houses sold in a given period rose, a simple average of transacted house prices would increase even if the price of both types of houses were unchanged.

Quality adjustment

Constant-quality measures of house prices try to standardise, and make comparable over time, the information available in the data, to overcome the limitations of simple averages. Three main methods are used for this: hedonic regression; mix adjustment; and the repeat-sales method.

Hedonic regression

The price of a house depends on its location and its physical characteristics. Hedonic regression is a way of estimating the value the market places on each of these attributes. For example, hedonic regressions can be used to estimate the market value of a bedroom, a garage, or a particular location. These estimates are then used to

construct the price of a synthetic house with a representative amount of each characteristic.

In terms of the previous example, a regression model could estimate the implicit market values of 'detachedness' relative to 'terracedness' in each period. If the model is correctly specified, the estimated prices of each characteristic, and therefore the estimated price of a house with fixed characteristics, will not be affected systematically by changes in the composition of the sample. (See the box on hedonic regressions for more detail on this method.)

Mix adjustment

Mix adjustment is an alternative approach to hedonic regression for removing the effect of changes in the characteristics of the sample (the 'mix'). House price observations are grouped into sets or 'cells' of observations on houses with similar location and physical attributes. In terms of the previous example, the sample could be divided into two cells—detached houses and terrace houses. In practice, the level of detail in existing mix-adjusted indices is much greater. For example, the mix-adjusted index produced by the Office of the Deputy Prime Minister (ODPM) contains over 300 cells, and the level of detail extends to a cell exclusively for second-hand, semi-detached houses with six rooms bought by first-time buyers in the North-East region.

Once all the data in the sample have been allocated to a cell, the mean prices in each cell are calculated. These mean values are then weighted (see below for a discussion of how the weighting method affects the interpretation of the average). The resulting weighted mean cell price is the 'mix-adjusted' price. A change in the composition of the sample will alter the number of observations in each cell. But if the cells are defined sufficiently precisely, so that all the elements of the cell have similar prices and price trends, then such compositional changes will not systematically affect the mix-adjusted house price.

Chart 1 shows the effect that mix adjustment has on the data in the ODPM house price index. The mix-adjusted index rose about one third faster than the simple average of the data over the 1980s, suggesting that the sample of transacted houses shifted towards relatively cheap houses over this period (possibly due to large sales of council houses). The mix-adjusted index rose more slowly than the simple average over the following

Hedonic regressions

The hedonic or characteristics approach to price measurement is based on the hypothesis that goods in themselves do not provide utility; instead utility is assumed to be derived from the properties or characteristics of goods (see Lancaster (1966)). By extension, the price of any good can be thought of as summarising the prices of its underlying characteristics. For houses, it follows that the price of a house will depend on the purchasers' valuation, or implicit price, of each characteristic and the quantity of each characteristic in the house (eg its size, location, number of bathrooms, bedrooms etc). In general, both qualitative (eg house location) and quantitative (eg floor size) characteristics will be relevant to the overall price. The implicit prices can be used to calculate the price of a house with a standard set of characteristics, which represent the type of house the price index is designed to track.


The implicit prices of characteristics cannot be observed, because the normal unit in which transactions take place is a complete house. Hedonic regressions overcome this by using a sample of house prices and the associated house characteristics to estimate the implicit market price of a unit of each characteristic. Estimation is complicated by two factors. First, there is no objective criterion that can be used to determine exactly which characteristics are

relevant to the overall price of a house. Second, the way in which a characteristic contributes to the house price is not always obvious. For instance, does each successive bedroom contribute equally to the house price, or does a third bedroom contribute less than the second or first bedrooms? Judgments on these issues will determine the specification of the hedonic regression.

The table below shows that the hedonic regressions employed in the construction of the Halifax and Nationwide house price indices differ to some extent in the judgments they embody: each is based on a somewhat different set of characteristics, and some characteristics contribute to the house price in different ways. For instance, both lenders assume that the number of bathrooms affects the price of a house. But the Halifax index treats each successive bathroom as contributing the same additional amount to the house price, whereas the Nationwide index makes no distinction between a house with two bathrooms and one with three or more.⁽¹⁾ Such discrepancies may give rise to differences in the two indices' estimates of the rate of house price inflation. In particular, the inclusion of a variable in one equation but not another is likely to affect the coefficients on other variables in the equation and therefore possibly the price of the typical house. It is difficult to determine how large these effects would

(1) The method used to construct the Halifax index is discussed in some detail in Fleming and Nellis (1984), available on request from the Halifax.

Chart 1
Ratio of simple average to mix-adjusted ODPM indices



five years, as the sample shifted back towards higher-value house types. So a simple average such as the Land Registry index would generally have understated constant-quality house price inflation over the 1980s, and overstated it over the 1990s.

Repeat sales

Both the hedonic regression and mix-adjustment approaches require a large number of dwelling characteristics to be recorded if they are to be reliable. In some cases this information is not readily available. Instead, there may be information on the history of transactions for a large sample of dwellings, which allows us to examine the price changes of individual houses.

Characteristic	In the Nationwide regression?	In the Halifax regression?
Detached house	✓	✓
Terrace house	✓	✓
Detached bungalow	✓	} ✓ uses one bungalow dummy variable rather than two
Semi-detached bungalow	✓	
Purpose-built flat/maisonette or new converted	✓	} ✓ uses one flat dummy variable rather than two
Converted flat/maisonette	✓	
Tenure	✓	✓
Number of bedrooms	✓	✗
Number of habitable rooms	✗	✓
Double garage	✓	✗
Number of garages	✗	✓
Number of garage spaces	✗	✓
Parking space or no garage	✓	✗
Central heating type	✓	✓
Floorsize (sq. ft.)	✓	✗
Number of acres	✗	✓
More than one bathroom	✓	✗
Number of bathrooms	✗	✓
Number of toilets	✗	✓
Garden	✗	✓
Subject to a road charge	✗	✓
Property age	✗	✓
New	✓	✗
Region	✓	✓
ACORN ^(a) classification	✓	✗
Parliamentary constituency	✓	✗

(a) A classification of residential neighbourhoods.

be, but even if the Halifax and Nationwide used the same data and definition of a typical house, their estimates of the price of a typical house would be likely to differ.

The Halifax and Nationwide indices are also constructed differently from the estimated prices of characteristics: the Halifax index uses a fixed

definition of a typical house, whereas the Nationwide index uses a definition that changes periodically, usually each year. The Halifax 'standard house' is defined by the characteristics of the average house on which the Halifax approved a mortgage in 1983. The Nationwide 'standard' house is, in general, defined by the characteristics of houses in rolling averages of the Survey of Mortgage Lenders, the Land Registry and Nationwide transactions. The Halifax index is, as a result, more sensitive than the Nationwide index to price movements in—for example—the North, because its weighting scheme places greater weight on price information from that region. Both indices are constructed from separate estimates for every period of the characteristics' prices, so the relative prices of characteristics are free to vary from one period to the next;⁽²⁾ and both draw on price data from houses on which these lenders have themselves granted mortgages.

In times of uniform market-wide inflation, differences in the standard house will probably only have a very small effect on measured inflation rates because all houses will be increasing in value at similar rates. But at times when inflation rates diverge between regions or house types, there can be some effect on measured inflation rates. The extent of the effect will depend on the sensitivity of each index to the various segments of the market.

(2) Both of the lenders' indices are based on separate hedonic regressions for each period's data. This is the least restrictive method that could be used, but it could cause what might be considered strange results: the price of a characteristic, say the price of a bathroom, could change sharply from one period to the next. An alternative approach would be to constrain the relative prices of characteristics to be constant for several periods, but allow them all to vary by a similar amount.

Observing the sale prices of a given house at two points in time will give an estimate of general house price inflation between these two transactions. With a sufficient number of estimates from partially overlapping periods, we can estimate, say, yearly house price inflation. For example, if one group of houses was sold in 2000 and again in 2002 at prices 15% higher, and another group suggests that prices rose 10% between 2001 and 2002, we can infer that prices rose by around 5% between 2000 and 2001. No repeat-sales indices yet exist for the United Kingdom, but they have been estimated on US data.

All these methods have disadvantages. Both hedonic regression and mix adjustment require that all the characteristics of the house that may affect its price are

controlled for. If some were omitted from the hedonic regression or cell structure, a change in the distribution of these characteristics over time would create inaccuracies in the estimated change in the price of a constant-quality house. For example, if fitted kitchens became more common, but were not recorded as a characteristic of the houses in the sample, the price index would rise too quickly: higher prices from the inclusion of fitted kitchens would be mistaken for an increase in the price of a constant-attributes house. To the extent that the existing house price indices do not measure such quality improvements as they become more prevalent, they would overstate the rate of constant-quality house price inflation. Furthermore, if these unobserved attributes were more common in properties sold at certain phases of the cycle (for

example, if the top end of the market were relatively active during booms) then the amplitude of fluctuations in house price inflation may be understated or overstated accordingly. In other words, hedonic and mix-adjusted indices will suffer from the same problems as simple averages if they do not control for all relevant, varying characteristics.

Repeat-sales indices are based on individual houses' inflation rates. Changes in the sample mix of price levels would not affect the estimate of inflation, but variation in the mix between houses with different inflation rates would affect the estimates of simple repeat-sales indices. For example, if detached and terrace houses appreciate at 0% and 5% per year respectively, a shift in the sample towards terrace houses will increase the estimated average inflation rate, because no account is taken of the characteristics of the sample. Furthermore, the estimated appreciation rate will also be biased if the property changes in condition or is altered between the two price observations. Hybrid hedonic repeat-sales indices remove this problem by controlling for the characteristics of the sample, or by treating past sales prices of the same house or nearby houses as control variables in hedonic regressions.

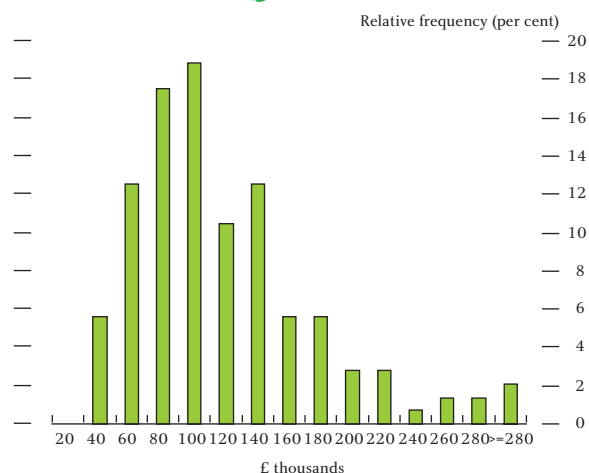
The meaning of 'average' house prices

Once the data have been standardised to remove the effect of variation in the sample, the question remains: which house or group of houses should the index represent? The levels and inflation rates of house prices in the United Kingdom are widely dispersed—the prices of similar dwellings in different locations can vary by a factor of more than seven (see Chart 2)—so this is a key question. There are two main issues.

First, what constant notional house or set of houses should the index try to represent? Should it be typical of the flow of transactions, or the stock of dwellings? The two will vary to the extent that different kinds of houses are transacted at different rates.

Second, should the more expensive houses in this set receive more weight, commensurate with their share of the expenditure on or value of the set? We can choose to represent the price of a house with typical characteristics, where all houses in a set have equal weight in determining what is typical: this is called 'transactions weighting' or 'volume weighting'. Alternatively, we can represent the price of a representative collection of houses, where more

Chart 2
Distribution of average transacted flat and maisonette prices in localities in England and Wales in 2002 Q3



Source: Land Registry.

expensive houses have an accordingly higher weight: this is 'expenditure weighting' or 'value weighting'. These two measures will in general behave differently.

To see this, suppose that there are equal numbers of detached and terrace houses in the housing stock. If the price of expensive detached houses were rising at 10% per year, whereas the price of cheap terrace houses were falling at the same rate, the inflation rate of the 'typical house', in which each type would have equal weight, would be zero. But the cost of purchasing the entire housing stock, or a representative share in it, would be rising: the extra cost of the detached house would outweigh the lower cost of the terrace house, even though the prices were changing at the same rate, but in opposite directions.

In other words, when the price index is constructed to reflect the value of the reference set (ie the value of the housing stock or typical housing transactions), rather than the value of a typical member of the reference set, expensive houses are given more weight. If all houses were appreciating at a common rate, both price indices reflecting value and volume weights would appreciate at this rate. But if low and high-value houses were to exhibit different price trends, the inflation rates of volume and value-weighted indices would diverge.


How do the existing measures of house prices fit into this framework?

There are several available measures of UK house prices. The four longest-established UK house price indices are produced by the Office of the Deputy Prime Minister

(ODPM), the Land Registry and the mortgage lenders Halifax and Nationwide. More recently, the housing web sites Hometrack and Rightmove have augmented this set. In addition to these indices, the Royal Institution of Chartered Surveyors and the House Builders Federation undertake qualitative surveys of changes in house prices.

Chart 3 shows that the four longest-established indices have similar long-term trends. The differences in their estimates of house price inflation will depend on the samples they are based on, and how these samples are transformed to produce the indices.

Chart 3
Halifax, Land Registry, Nationwide and ODPM indices of house prices^(a)



(a) Nationwide (pre-1991) and Land Registry indices seasonally adjusted by the Bank of England.

Table A shows that, while some indices are weighted to represent a flow of transactions, others represent a stock of dwellings. This will cause the indices to diverge to the extent that the average transacted house is different from the average house in the stock, and the price trends between them vary.

Table A
UK house price indices

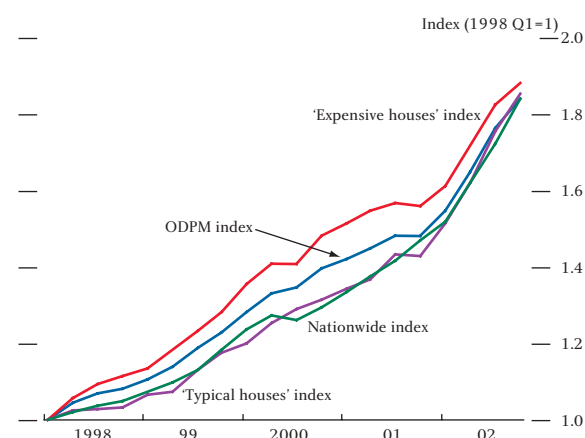
Index	Sample	Standardisation method	Seasonally adjusted?	Weights used	Weighting method
ODPM (a)	5% sample of Council of Mortgage Lenders' eligible completions	Mix adjustment	No	Rolling average of UK transactions	Expenditure
Halifax	Loans approved for house purchase	Hedonic regression	Yes	1983 Halifax loan approvals	Volume
Nationwide	Loans approved for house purchase	Hedonic regression	Yes	Rolling average of Survey of Mortgage Lenders, Land Registry and Nationwide transactions	Volume
Land Registry	100% of sales registered in England and Wales	Simple average	No	None	Expenditure
Hometrack	Survey of approx. 4,000 estate agents' estimated local average prices	Mix adjustment	No	England and Wales housing stock	Expenditure
Rightmove	Sellers' asking prices posted on web site	Mix adjustment	No	England and Wales housing stock	Expenditure

(a) The ODPM is in the process of expanding the Survey of Mortgage Lenders data set on which its index is based, and will shortly switch to a variant of the hedonic regression method.

Both the Halifax and Nationwide indices represent the price of a typically transacted house—they are 'volume-weighted' indices of typically transacted house prices. The Halifax typical house has the mean characteristics of the houses on which this lender approved mortgages in 1983, whereas the Nationwide index weights are derived from rolling averages of the Survey of Mortgage Lenders, Land Registry and Nationwide samples. In contrast, the ODPM index is 'value weighted'—it is effectively a deflator for typical housing market transactions.

To demonstrate the practical effect of the difference between volume and value weights, Chart 4 shows two subindices based on the most expensive quartile and the middle-priced 50% of the cells in the ODPM index, which can be interpreted as indices of 'expensive' and 'typical' houses. As expected, the volume-weighted Nationwide index seems to track the typical house

Chart 4
Levels of 'expensive' and 'typical' house price indices, the ODPM index and the Nationwide index



index more closely than does the value-weighted ODPM index, which moves more in line with the ‘expensive’ index.

The Land Registry index is based on a 100% sample of housing transactions in England and Wales, so it should not be subject to sampling error.⁽¹⁾ But, as Table A shows, the index is not mix adjusted, so variations in the mix of transactions affect the level of the index. Chart 5 shows that a seasonally adjusted, partly mix-adjusted version of the Land Registry index is smoother than a seasonally adjusted version of the published simple average. Had the latter been used as a guide to the price of a constant-quality house rather than as the average price of transactions, it would have given a misleading picture on many occasions. For example, the seasonally adjusted inflation rate of the headline index fell in 2001 Q4, while that of the mix-adjusted version rose.

Chart 5
Simple average and mix-adjusted Land Registry indices


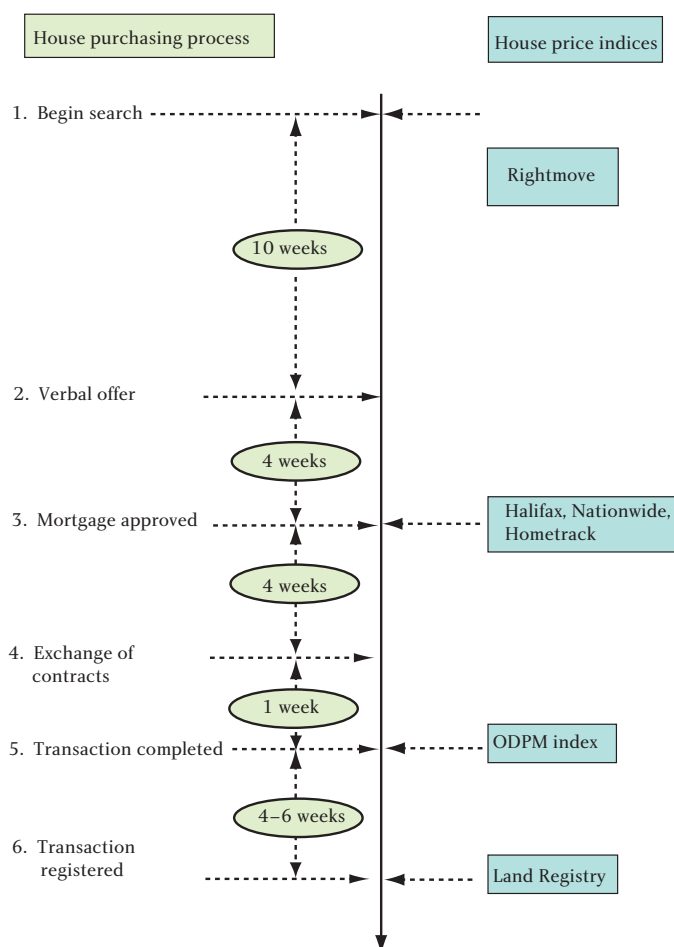


Chart 6 shows a timeline for a typical house purchase. The indices in Table A measure house prices at different points along this timeline, so indices nearer the

Chart 6
House purchase timeline and house price indices



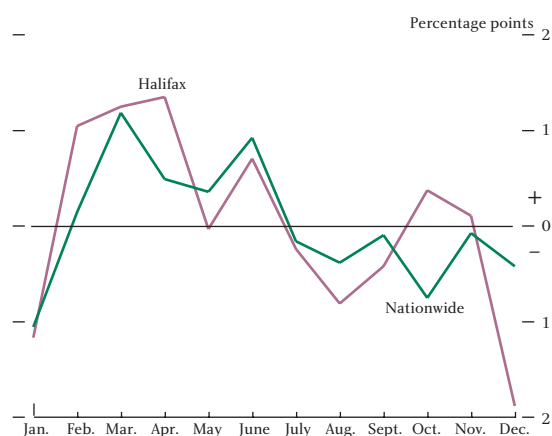
Sources: Bank of England and ODPM.

(1) Note that the latest observation of the headline Land Registry index is usually revised upwards in the following release. For example, the 2002 Q3 observation was revised upwards when the 2002 Q4 data were released. The requirement to pay Stamp Duty, from which lower-value properties are exempt, delays the arrival of some observations in the sample. So data arriving late in the sample tend to have above-average prices, causing the upward revision.

beginning may be likely to detect a change in house prices first: the house prices appearing in the January Halifax index will not appear in the ODPM index until February. But the lagged comovement of the indices will be weakened to the extent that the sale price of a given house changes throughout the process. For example, prices may be renegotiated between mortgage approval and completion stage.


Aside from issues of measurement, a further complication is that not all of the indices are seasonally adjusted. The seasonal component of house price inflation is generally positive from February to July and negative in the rest of the year, and it can change sharply from month to month (see Chart 7). For example, the seasonal components of monthly Halifax and Nationwide inflation increase on average by 1.7 percentage points in February. So we should expect large increases in the monthly inflation rates of seasonally unadjusted indices in February, even when 'underlying' inflation is flat.

Chart 7
Seasonal factors in monthly Halifax and Nationwide inflation



Finally, house price indices remain volatile even after they have been adjusted for seasonality and the mix of transactions. This does not necessarily reflect noise or

Chart 8
Changes in 'trend' and 'actual' monthly Nationwide inflation^(a)



(a) Centred moving-average trend given by X-12 filter.

error in the indices: the 'true' variable they are measuring may also be volatile. However, policy-makers and commentators are frequently interested in the medium-term trend of house price inflation. As Chart 8 shows, most of the changes in monthly inflation each month do not reflect changes in an estimate of this trend, but are instead frequently reversed. This means that monthly changes in house price inflation do not individually contain much information about whether the medium-term inflation rate of house prices is rising or falling: several months of data are usually necessary to establish a change in the trend.

Conclusion

The heterogeneity of housing makes constant-quality prices difficult to measure, and means there are several valid concepts of 'house prices', each of which can behave differently. The indices available for the United Kingdom embody a variety of these concepts. Observers and policy-makers must be careful to match the measure they use with the concept they are interested in, and to ensure that the information in short-run changes in house price inflation is not overstated.

References

Fleming, M and Nellis, J (1984), *The Halifax house price index: technical details*.

Goolsby, W C (1997), 'Assessment error in the valuation of owner-occupied housing', *Journal of Real Estate Research*, Vol. 13:1, pages 33–46.

Lancaster, K (1966), 'A new approach to consumer theory', *The Journal of Political Economy*, Vol. 74(2), pages 132–57.