Consumer Price Indices in the UK

Main Findings

The report *Consumer Price Indices in the UK,* written by Mark Courtney, assesses the array of official inflation indices in terms of their suitability as an uprating index that measures the purchasing power of wages and pensions. The main findings of this report are set out below:

- Overall, taking account of both coverage and formula effect differences and within the limitations of how price data is collected within the UK, the Retail Prices Index (RPI) is as good a consumer price index as one can get for uprating purposes. The systemic differences between the RPI and the Consumer Prices Index (CPI) are the result entirely of under-estimation by the CPI. This under-estimate is currently (2014) running at 0.9 percentage points per annum.
- The coverage of the RPI is targeted on the working population, as it excludes pensioners, tourists and the wealthiest 4% of households. In contrast to CPI, it also includes the owner occupier housing costs that form a major element of most household expenditure. These differences in coverage have caused the CPI to be lower than the RPI by an average of 0.3 percentage points per annum over the seventeen years since the CPI was introduced.
- The RPI and CPI use different statistical methods (known as aggregation formulae) for calculating average price changes for many of the items covered. The RPI uses an arithmetic average, usually in the form of the Average of Relatives, whereas the CPI, following the requirements of the EU's Harmonised Index of Consumer Prices, uses the Geometric Mean. Because the UK has unusually broad definitions of its items, with consequent variations of inflation within items, this formula effect difference is unusually large in the UK: it has averaged 0.6 percentage points a year since the introduction of the CPI. The choice of aggregation formula is therefore the most important difference between the RPI and CPI.
- The "stochastic approach" used by statisticians as a method of assessing price indices shows that the Average of Relatives used by the RPI is an unbiased estimator of the average rate of inflation, whereas the Geometric Mean used by the CPI is biased downwards. This theoretical result has been supported by empirical research published by the Office for National Statistics (ONS). This means that the formula effect represents an additional under-estimate of inflation by the CPI by, on average, 0.6 percentage points.
- This under-estimate when using the Geometric Mean has been defended in the past on the grounds that it is allowing for consumer substitution towards goods whose prices have risen more slowly. But this neglects the fact that price changes are driven by changes in both supply and demand, not in supply alone, so that trying to allow for consumer substitution does not favour one aggregation formula over another. The ONS now accepts that this Economic argument is no reason to modify the usual stochastic approach.

- It is sometimes argued that the process of re-weighting the RPI and CPI every year to allow for changing patterns of expenditure will introduce an upward bias, due to "price-bouncing," in the RPI. However, there is no theoretical reason why any bias should be upwards and the available empirical evidence is that any such effect is very small – much less than 0.1 percentage points per annum.
- Although the Average of Relatives is widely used internationally in the calculation of other price indices, such as producer price indices or trade price indices, the UK is almost alone in using it in a consumer price index, and it has been alleged that it is therefore not consistent with best international practice. In fact, few countries have ever used the Average of Relatives for their consumer price index, since many of them have tightly defined, homogeneous items, unlike the UK, and use the Ratio of Averages (a different formulation of the arithmetic average) throughout. Fifteen years ago Australia and the USA switched their consumer price indices from using the Average of Relatives to the Geometric Mean and their example might be relevant, except that their reason for switching was to account for consumer substitution, and the ONS now considers that the Economic approach that gave rise to the idea of consumer substitution does not provide even a weak reason for such a switch.
- In March 2013, the UK Statistics Authority downgraded the classification of the RPI from National Statistic to official statistic. It made clear that it was doing so solely on the basis of evidence presented by the ONS, without following its usual practice of written consultation or discussion with outside experts, and the substantive reason for its decision was that the National Statistician had decided that the ONS would henceforth perform only routine updating of the RPI, which violated the UK Statistics Authority's requirement for continuous improvement. In the circumstances, the withdrawal of National Statistic designation from the RPI lacked a convincing statistical basis and cannot be regarded as a comment on its current accuracy.
- Attempts have been made to add owner-occupier housing expenditure to the CPI by introducing a new, experimental price index, CPIH. This uses a rental equivalence measure for housing costs which is a very poor proxy for expenditure that home owners actually incur and has been accompanied by data problems. Consequently, CPIH has so far gained little acceptance.
- There is also a new price index RPIJ, which uses the Geometric Mean within the RPI, wherever it previously used the Average of Relatives. This effectively turns it into a CPI with an RPI-type measure of owner-occupier housing expenditure (and some other coverage changes). Since it still has the under-estimate of inflation inherent in its use of the Geometric Mean it is unlikely to gain acceptance as an uprating index.

Consumer Price Indices in the UK: Executive Summary

It is generally accepted that there are advantages to having more than one consumer price index, in view of the different uses to which such an index can be put. The two most important uses are as an indicator for monetary policy and, as an uprating index, to measure the purchasing power of wages and pensions. Although the paper refers to the properties of a good monetary policy indicator, our main concern is with the second type of index.

Even though there is scope for more than one consumer price index, the UK has ended up with a confusingly large number – nine at the last count that have some claim to apply to the population at large, as well as two or three more that aim to measure the prices faced by sub-groups within the population. However, our main focus is on just the two most prominent ones, from one or other of which all the rest derive. These are the Retail Prices Index (RPI), which was developed as an uprating index but, in a slightly modified form, was also used as a monetary policy indicator from 1975 until 2003; and the Consumer Prices Index (CPI), the name given in the UK to the EU's Harmonised Index of Consumer Prices (HICP), which was developed as a monetary policy index but has also been used as an uprating index in the UK since 2011.

Both the RPI and CPI are, in principle, *fixed-basket* indices. They aim to establish the cost of a representative basket of consumer goods and services in the base period and then to compare that with the cost of the same basket of goods and services in subsequent periods. Ideally, the basket would contain as wide as possible a range of goods and services: in practice, about 700 items are included, either because they are important in themselves or because they are taken as being representative of a wider class of goods.

The price indices calculated for the representative items are aggregated, by taking a weighted arithmetic average (a *Laspeyres-type index*) on the basis of the previous year's Living Costs and Food Survey of household expenditure, to obtain the price indices for the 85 sections of the RPI (or 85 classes of the CPI). These are then aggregated further into the overall RPI or CPI.

Since, with one or two exceptions, the RPI and CPI are calculated on the basis of the same sample of prices (about 120,000 prices collected each month) and by the same methods, why are they so different, with the CPI inflation, on average, having been 0.8 percentage points less than RPI inflation over the period 1997-2013? There are two reasons: differences in coverage and differences in aggregation formula.

The RPI was developed as an uprating index, and therefore aims to include all expenditure by UK households, including on foreign holidays, but excluding spending by the richest 4% of households (who account for 10.5% of household expenditure) and by pensioners mainly dependent on state benefits (who account for 2.2% of household expenditure). The CPI, developed as a monetary indicator, aims to include all monetary expenditure within the UK, including by foreigners, but the main difference in coverage is that the CPI, because it is constrained by the regulations governing the HICP, excludes almost all owner-occupied housing costs. Because housing costs have generally risen faster than other inflation,

differences in coverage mean that the CPI has, on average, been 0.3 percentage points less than RPI inflation.

If differences in coverage were the only source of divergence, the choice between RPI and CPI would be clear: the RPI is better as an uprating index, the CPI as a monetary policy indicator. However, there is also a formula effect difference, due to the different ways the RPI and CPI are calculated, which averaged 0.5 percentage points up to 2009, and has averaged 0.9 percentage points thereafter.

We have seen that, in principle, the RPI and CPI are calculated as fixed-basket indices in the same way, with the inflation rates of representative items being aggregated by taking an arithmetic average, weighted by the item's importance in household expenditure. However, when it comes to calculating the inflation rate of the 700 representative items that go into the index, we usually have no data on the expenditure associated with each of the price quotes that are used to estimate the inflation rate for each item (some of the items are in turn stratified by region or type of outlet, so there are actually more than 700 "elementary aggregates"). Thus, if our representative item is "womens' formal skirt," and we have price quotes from 20 widely dispersed shops with quotes for a skirt in the base month and the current month – different skirts in each shop, but the same skirt in a particular shop each month – we don't know what the sales in each month are. We therefore have to use an unweighted average of the 20 sampled prices.

To calculate this unweighted average, the RPI uses one of two forms of the arithmetic average: it uses the *Ratio of Aver*ages (sometimes called the Dutot) when the items are homogeneous, and the *Average of Relatives* (sometimes called the Carli) when they are not. The weighted Average of Relatives is the formula that is used by both RPI and CPI for upper-level aggregation, so the RPI is *consistent in aggregation*. The CPI generally uses the *Geometric Mean* (sometimes called the Jevons) – a geometric mean of n numbers is calculated by multiplying them all together and taking the n^{th} root. As a matter of arithmetic, a geometric average is always lower than an arithmetic average, so the Geometric Mean is always less than the Average of Relatives (with the Ratio of Averages it can go either way, but the formula difference is caused almost entirely by those items where the RPI uses the Average of Relatives).

Which of these is correct? If we take the usual statistical or *stochastic approach*, we regard each of our matched price quotes as giving rise to a price relative (the price in the current period divided by the price in the base period) which is an estimate of the price relative for the item as a whole. In the absence of any knowledge about the importance of the price quotes, the usual statistical assumption is that they are equally important and distributed randomly around the overall item price relative. Then it is an established result that the Average of Relatives will be an unbiased estimator of the item price relative (i.e. the item inflation rate plus 1), whereas the Geometric Mean will be biased downwards. And this is the case regardless of how price relatives evolve or what statistical distribution they have.

However, perhaps we are not totally ignorant about the expenditure associated with each price quote. We might draw on Economic theory to give us an assumed relationship between price and quantity, as in the *Economic approach*. Here, we start by observing that the mix of purchases within an item in the current period is not, for one reason or another, exactly the same as in the initial period. We might want to take account of this by giving an equal weight to expenditure in the current period as to expenditure in the base period (we will than, strictly speaking, no longer have a fixed-basket index). In this case we are led to a symmetric, weighted price index, such as the *Fisher index*. If we can assume that expenditure on each sampled product is the same in both the current and base period (i.e. price and quantity changes are exactly offsetting), then the Geometric Mean is a good estimator of the Fisher index, and the Average of Relatives overestimates it. This assumption would be justified if the prices we observe represented movements along a fixed demand curve and there was a high level of consumer substitution. Conversely, if there was no or little consumer substitution, we could assume that quantities in the current and base periods were approximately equal, and the Average of Relatives would be a good estimator of the Fisher index and the Geometric mean would underestimate it. This is the basis on which, in 2011, the ONS represented the choice of appropriate aggregation formula as being dependent on the degree of consumer price substitution.

But this is far too narrow an approach. It is an elementary proposition of Economics that price and quantity movements are caused by changes in supply *and* demand. The price changes we observe are almost never caused by changes in supply alone, such as would let us observe movements along a static demand curve. The opposite case would be a price change caused just by a shift in demand, letting us observe movements along a fixed supply curve. In that case, it turns out, the Average of Relatives will always underestimate the Fisher index and the Geometric Mean will underestimate it by even more. In the usual case where there are both supply and demand shocks, there is no clear outcome. However, we know that the Geometric Mean will always be biased down for demand shocks and will be biased down for supply shocks if consumer substitution is low, whereas the upward and downward biases of the Average of Relatives are more evenly distributed, so this is a weak indication that the Average of Relatives would be a better approximation of the weighted Fisher index.

The stochastic approach quite strongly and the Economic approach more weakly point to a downward bias by the Geometric Mean. However, this needs to be tested empirically, since we cannot know from theory how well the unweighted elementary aggregates will estimate our target weighted indices. The ONS research team conducted a large-scale study with matched price and quantity data from the UK alcohol sector. Where the target was the Laspeyres index, which is the index used in upper level aggregation, the results were clear-cut: the Average of Relatives performed remarkably well as an estimator of the Laspeyres index. Where the target was the Fisher index, the Average of Relatives and Geometric Mean performed equally well: the ONS team concluded that in this case the result was data-dependent and that it would be difficult for an unweighted index to be a good estimator for that type of weighted index (a result confirmed in their other research paper, where none of

the unweighted indices was a good estimator of a population Fisher index, although the Average of Relatives and Ratio of Averages were better than the Geometric Mean).

There is one final consideration. Both the RPI and CPI are *annually chained* indices, with the expenditure weight revised each year to take account of new goods or shifts in expenditure patterns. This is usually considered a good thing, but it means there is some *chain drift*, where the chained index differs slightly from what the direct index would be, calculated on unchanged weights. Since the Average of Relative is not a *transitive* index, it adds some additional chain drift, one way or the other: where inflation is relatively smooth, it will have a negative effect, and where prices oscillate, and one gets price bouncing, it will have a positive effect. Where, unusually, price indices are chained monthly (as in Canada and Denmark) this price bouncing effect can be strong for goods with a seasonal price pattern or with special pricing strategies. But for annual chaining it is always much weaker. The best, indirect evidence we have for the UK is that use of the Average of Relatives has a net positive effect on chain drift, but only by about 0.02 percentage points. And the ONS has said that it cannot provide a numerical estimate of price bouncing, since the effect would be indistinguishable within the general chain drift that affects both RPI and CPI. One must therefore conclude that annual chaining explains effectively none of the formula effect difference between RPI and CPI.

Overall, taking account of both coverage and formula effect differences and within the limitations of how price data is collected within the UK, the RPI is as good a consumer price index as one can get for uprating purposes. The systemic differences between the RPI and the CPI are the result entirely of under-estimation by the CPI.

Three additional issues:

1) Although the Average of Relatives is widely used internationally in the calculation of other price indices, such as producer price indices or trade price indices, the UK is almost alone in using it in a consumer price index, and it has been alleged that it is therefore not consistent with best international practice. In fact, few countries have ever used the Average of Relatives for their consumer price index, since many of them, particularly in Europe and Asia, have tightly defined, homogeneous items, unlike the UK, and use the Ratio of Averages throughout. Fifteen years ago Australia and the USA switched their consumer price indices from using the Average of Relatives to the Geometric Mean and their example might be relevant, except that their reason for switching was to account for consumer substitution, and the ONS now considers that the Economic approach does not provide even a weak reason for such a switch.

2) In March 2013, the UK Statistics Authority downgraded the classification of the RPI from National Statistic to official statistic. It made clear that it was doing so solely on the basis of evidence presented by the ONS, without following its usual practice of either written consultation or discussion with outside experts, and the substantive reason for its decision was that the National Statistician had decided that the ONS would henceforth perform only routine updating of the RPI, which violated the UK Statistics Authority's requirement for

continuous improvement. In the circumstances, the withdrawal of National Statistic designation from the RPI lacked a convincing statistical basis and cannot be regarded as a comment on its current accuracy.

3) Attempts have been made to include owner-occupier housing expenditure within the CPI. There is an experimental series, CPIH, which uses a rental equivalence measure. This has data problems and is a very poor proxy for expenditure that home owners actually incur, and has so far gained little acceptance.

There is also a new series RPIJ, which uses the Geometric Mean within the RPI, wherever it previously used the Average of Relatives. This effectively turns it into a CPI with an RPI-type measure of owner-occupier housing expenditure (and some other coverage changes). Since it still has the under-estimate of inflation inherent in its use of the Geometric Mean it is unlikely to gain acceptance as an uprating index.

This report was written by Dr Mark Courtney, formerly Deputy Director and Head of Economics in the Regulatory Impact Unit, Cabinet Office. He previously worked as an Economic Adviser in HM Treasury and as Senior Lecturer in Economics at Rhodes University. He was nominated by the Royal Statistical Society in 2011 to serve on the technical advisory group to the Office of National Statistics on clothing inflation, and has published a critique of UK consumer price indices in the *Statistical Journal of the International Association of Official Statisticians*. The need for a paper to bring clarity to the proliferation of consumer price indices was identified by public sector union UNISON, who also provided a research grant. Responsibility for the analysis and opinions in this report is the author's alone.