General Insurance Seminar

Connecting Today and Tomorrow

13 – 15 November 2016
Grand Hyatt • Melbourne

Actuaries Institute

The rise and fall of inflation indices

Prepared by Hugh Miller

Presented to the Actuaries Institute
General Insurance Seminar
13 – 15 November 2016
Melbourne

Institute of Actuaries of Australia
ABN 69 000 423 656
Level 2, 50 Carrington Street, Sydney NSW Australia 2000
t +61 (0) 2 9239 6100 **f** +61 (0) 2 9239 6170
e actuaries@actuaries.asn.au **w** www.actuaries.asn.au

Abstract

The paper is a curiosity-driven investigation into inflation indices, with a particular emphasis on the consumer price index (CPI) and average weekly earnings (AWE) index. We explore some of the detail in how these indices are defined and constructed, and then discuss some potential implications. The main conclusions are:

- Inflation indices are highly definitional and reasonably complex to construct. These characteristics mean that they are not always the most appropriate index to apply, and perhaps we should have more confidence in altering the index to suit a particular need.
- Data collected and reported by the Australian Bureau of Statistics (ABS) is rich, despite budget pressures and limitations on collection frequency. A knowledge of what is available helps broaden understanding. For example, the subcomponents reported along with CPI are very useful.
- Technology and innovation pose both a challenge and opportunity for robust data collection. Inflation indices are slow to reflect the digital economy and its role in prices.
- Inflation-linked bonds are more useful as a market estimate of future inflation now than they were five years ago, thanks to a marked increase in issuance.
- By focusing on the average, AWE inflation is sensitive to inequalities in wage growth. The median wage historically has grown more slowly, and in some contexts may be a more appropriate index of change.
- There are some significant considerations when moving between individual gross wages and household income. The differences may be important for some financial service contexts.
- The switch to half-yearly AWE reporting by the ABS causes difficulties for some types of actuarial work. Work-arounds should be considered, including the use of more heavily smoothed AWE indices.

TABLE OF CONTENTS

1 Introduction 3
1.1 Background 3
1.2 Aims of this paper 3
1.3 Other references 3
2 Understanding CPI 4
2.1 What does CPI measure? 4
2.2 How is it measured? 5
2.3 What is not measured? 6
2.4 Geometric means for products 7
2.5 Why is CPI hard to measure, and is it getting harder to measure? 7
2.6 Main biases in CPI 9
2.7 Future improvements to the CPI 9
3 Implications of the CPI for users 11
3.1 Choice of index 11
3.2 Sub-components of the CPI 11
3.3 Allowance for biases 12
3.4 RBA target band 13
3.5 Inflation-linked bonds 14
4 Understanding AWE 17
4.1 What does AWE measure? 17
4.2 What is not measured? 18
5 Implications of AWE indices for users 19
5.1 Why don't we use the median wage? 19
5.2 Income inequality in Australia 20
5.3 Is household wage or individual more relevant? 22
5.4 The switch to half-yearly AWE 22
6 Conclusions 25
7 References 26

1 INTRODUCTION

1.1 Background

Inflation is a fundamental component of the economy and insurance. It is the most direct goal of monetary policy, and is the primary source of claims increase over the long term. Many financial instruments and payment are tied to specific inflation indices.

Inflation is also a very complex and at times vague concept. The Australian Bureau of Statistics (ABS) produces many different indexes, with two of the most prominent being consumer price inflation (CPI) and average weekly earnings (AWE). CPI itself is complex in its design and reflects a number of different underlying trends in the economy. AWE can be volatile and is sensitive to factors such as wage inequality and the relative strength of different industries.

This suggests that there is some value in a detailed understanding of inflation and index creation.

1.2 Aims of this paper

This paper is primarily curiosity driven – it explores a number of different issues around inflation indices. Because of this, it may be of less direct relevance to general insurance. There are a number of useful existing papers that address specific topics, including:

- Inflation risk in general insurance, (De Ravin and Fowlds, 2010)
- Inflation forecasting (Miller, 2010)
- Superimposed inflation (Pearson and Beynon, 2007)
- Extreme events in inflation (Ahlgrim and D'arcy, 2012)

Despite the limited aims of this paper, it should give a reader more understanding and confidence in using inflation indices. I hope this will prove to be of some value to an interested reader working in insurance.

1.3 Other references

Most of the background material on CPI and AWE are drawn from data and papers from the ABS. We typically have not referenced the ABS throughout the paper, so as to not double its length. Interested readers are referred to the ABS website. We have included appropriate ABS references at the end of this paper, along with other papers referred to throughout.

2 UNDERSTANDING CPI

2.1 What does CPI measure?

The two main approaches to measuring inflation are a cost of living index (which attempts to adjust for spending patterns) or a 'basket of goods' approach (which holds spending categories fixed). The CPI is the second type, which is far easier to measure robustly over time. The basket is a weighted average of consumer spending categories, with the weights updated relatively infrequently (historically every five or six years).

There are 11 major categories of expenditure:

- Food and non-alcoholic beverages
- Alcohol and tobacco
- Clothing and footwear
- Housing
- Furnishings, household equipment and services
- Health
- Transport
- Communication
- Recreation and culture
- Education
- Insurance and financial services.

These categories are further subdivided and there are 87 subcategories ('expenditure classes') which are individually estimated and then recombined using weights. There are about 1,000 'elementary aggregates' (e.g. 'white sandwich bread') that feed the 87 expenditure classes. There is also lots of judgment and use of other data in choosing the 1,000 items (e.g. bread rolls are not used as not sold by weight, compared to bread loaves).

CPI is monitored for the 8 capital cities, with weights determined by the Household expenditure survey (HES). Current weights (the 16th) are from 2011, based on data collected over 2009-10. We are about due for another refresh in the next 18 months. Combining HES with the CPI index allows us to estimate volume changes over time too.

To give a concrete illustration, three of the 87 subcategories relate to footwear. The annual spend on each group is calculated from the HES. They are then converted to weights of 0.14%, 0.34% and 0.13% for men, women and children respectively.

Figure 1 Annual spending on footwear, 2009-10 household expenditure survey

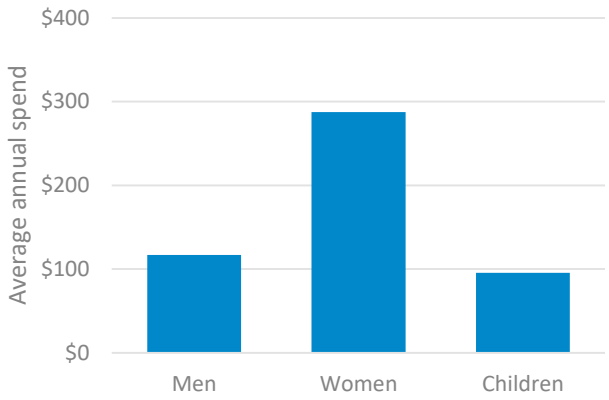

CPI is released quarterly, with release dates about a month after the end of a quarter. Prices are measured over the course of the quarter, so can be regarded as an 'average over quarter' level, rather than an end of quarter estimate.

2.2 How is it measured?

A schema (specific products and stores to check) is developed for each of the 1,000 elementary aggregates. A basket of representative products (with their own weighting within the schema) is developed. Field agents then collect those prices regularly in each of the eight capital cities, totalling about 100,000 prices collected per quarter.

Prices are generally collected 'in-store', although a number of government agencies and companies provide information directly. Electricity, gas and insurance prices are collected in this way which removes some of the complexity associated with individualised pricing in these industries.

The price volatility of an item will affect its collection frequency. Petrol, grocery items and airfares are particularly prominent for their price volatility.

CPI is a **quality-adjusted** index. So if a product is adjusted across time-periods, then a judgemental change will be applied to allow for it. This process is easier for some cases than others:

- A change in packaging size (200g package → 180g) can be translated into a simple price increase.
- A change to ingredients or other inputs requires some judgement as to quality change. Even packaging changes can be viewed as a change to quality.
- Such adjustments sometimes receive criticism since, if many products are 'improved' in quality in this way, it can lead to an impression of price deflation, even if consumer spending in the category remains the same.
- A change in channel is also a difference in quality; milk from a convenience store is different to the same milk in a supermarket.
- Audio-visual and computer prices are notoriously hard to monitor in light of quality improvements, due to rapid improvements in technology. The ABS matches models across months to measure price deflation due to such improvement, although this likely undervalues the class' influence over time (as people typically buy the improved model rather than spending less).

There are many other areas where judgement is required:

- Telecommunication prices are difficult due to bundling and non-comparable pricing (e.g. including a mobile phone in a plan, call costs). Other technological differences are quite difficult (e.g. quality difference between VOIP calling and standard fixed line).

» The choice of stores is clearly significant. Larger stores are typically easier to monitor, but will price differently to smaller stores. The relative market share of larger and smaller stores might be hard to estimate easily.

2.3 What is not measured?

The CPI does **not measure cost-of-living**, and very few statistical agencies globally have credible cost-of-living indices. A cost-of-living index, in theory, would track how consumer spending was changing between periods (both in absolute levels and the relative spending across categories). So a spike in banana prices would not be as dramatic in a cost-of-living index as people switch to other fruits. That said, the ABS maintains two different indices targeted towards cost-of-living: Analytical Living Cost Indexes for Selected Australian Household Types (cat. no. 6463.0) and the Pensioner and Beneficiary Living Cost Index (cat. no. 6467.0).

Also, it is useful to remember that CPI does not represent the costs for an 'average' household; few households would buy everything in the index in a quarter, and certainly not according to the specified weights. For example, rents and new house purchases have similar contributions to the index, but households usually do not pay both.

Third, CPI only relates to capital cities. There is no price visibility for country towns and rural areas, which in theory could have different price dynamics. This is perhaps most obvious for housing costs (see Section 3.3).

Fourth, there are some key spending class exclusions. These exclusions are made for a combination of conceptual and practical reasons:

» **Life insurance and superannuation costs:** Life insurance premiums are roughly 3% of GDP, so are clearly a significant part of consumer spending, but are difficult to measure effectively (on a premium minus benefits basis). Similarly, superannuation costs have historically not been included.
» **Non-life insurance:** Is included, but price changes are based on observable gross premiums. A more correct basis would be premiums less the value returned to consumers as claims, but this is hard to estimate.
» **Purchases of established houses:** Purchases of new houses (excluding land value) and spending on major renovations is included, but purchase of existing houses is not. This is consistent with the ABS's acquisitions approach to CPI measurement, but it may surprise some users that the majority of mortgage repayments are excluded from the index.
» **Online retailing:** Online-only retailers have traditionally been excluded from the indices, which will increasingly skew results. In particular, the impact of the app economy is completely invisible to CPI.
» **Indirect banking costs (particularly bank interest):** This is considered a borderline category by the ABS, but a big one; with roughly $1tr in outstanding mortgages, interest repayments are a large portion of household costs that are excluded. However, many people wouldn't necessarily think of a mortgage rate increase as 'inflation'.

In all, financial services are one of the most poorly measured categories in CPI.

Top level compositional changes are only allowed for slowly, via the irregular weight updates. This includes spending changes related to:

» Demographic changes (increased prominence of aged care, say)
» Technology (for example, electronic news and music)
» Consumer tastes and priorities (for example, changing mix of domestic and international travel).

Lower level compositional changes are sometimes updated more regularly. For instance, a switch to artisan breads might increase their weight in the bread category. This reweighting does not change the price level – just that future artisan bread price increases will have more consequence in CPI changes.

Finally, CPI is a relative index, so caution is needed in comparing the absolute levels of the index. In particular, if the Adelaide CPI level is below the Perth CPI level, this does not mean that prices are lower in Adelaide. Instead, you can only really compare how price changes have differed over a specified time period.

2.4 Geometric means for products

Each of the 1,000 elementary aggregates has its own basket (with sub-weights). There has been significant investigation on the 'right' way to combine price increases to get the change for the elementary aggregate. Options include:

- Calculating the average price in quarter 1 and quarter 2, then taking the difference.
- Calculating the change in price for each item and then averaging (arithmetic mean) change across them
- Calculating the change in price for each item and then averaging (geometric mean) change across them.

The third approach (geometric mean, also called the 'Jevons index') is generally used. As the geometric mean is always less than or equal to the arithmetic mean, it is viewed as a way to partly allow for product substitution within the elementary aggregate (a 'constant elasticity of substitution index'). It is equivalent to assuming that the total spend on any individual product is constant (so a 1% increase in price will be met with a 1% decrease in spend on that product).

There are exceptions to the geometric mean too. In situations where there is clearly no elasticity, such as council fees, an arithmetic mean is used. Also, in cases where a price falls to zero (or has a very large price reduction) the geometric average can be very low, so is sometimes adjusted.

2.5 Why is CPI hard to measure, and is it getting harder to measure?

In some ways CPI is getting easier to measure. For example, scanner data is increasingly being fed to the ABS, which allows for more automated and accurate price collection. However, in a number of important ways CPI, as well as GDP, is getting harder to measure.

2.5.1 Changing weights and the transition to services

While physical goods, like food and clothing, are tangible qualities, consumer spending has moved away from these and towards services.

Figure 2 Change in CPI category weights from 1982 to 2011

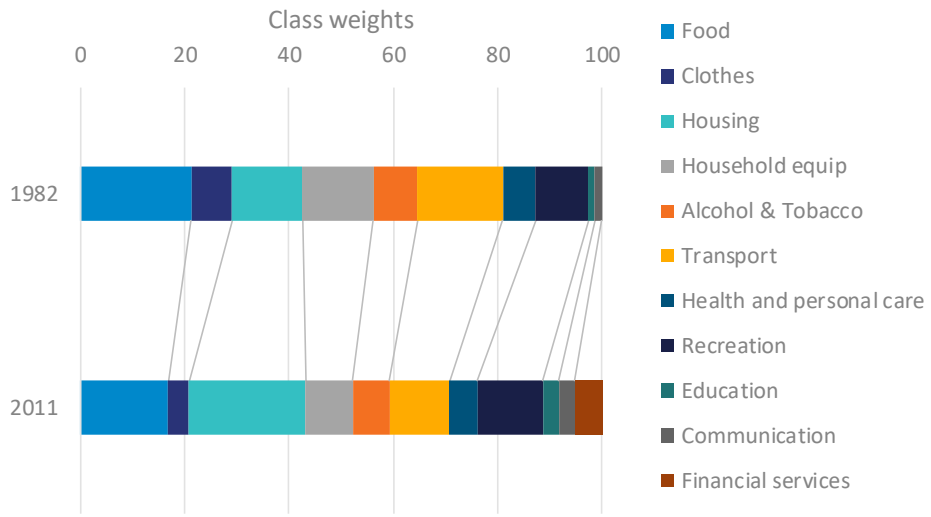

Housing has grown rapidly, from 13.5% to 22.3%. Most physical goods have fallen (in relative weight, not necessarily absolute spending). For example, clothing has halved to about 4%. By contrast services have increased. A new category for financial services was added which is now 5%, the communication category has doubled (even though app and streaming spending is not yet counted) and education has increased substantially too.

Services are harder to measure. Conceptually there is a distinction between:

- Length of service provided (e.g. price per hour of lawyer fees)
- Outcome based pricing (e.g. price per conveyancing contract).

The ABS favours the latter, but there are still significant issues. Classification is much less complete compared to manufactured goods. Quality is very hard to gauge, and the variation in pricing (due to quality or otherwise) tends to be large. The categorisation and structure of services is more difficult too; although services are now two-thirds of GDP, there are many fewer subcategories for services compared to more concrete sectors in agriculture and manufacturing. GDP makes use of the Australian and New Zealand Standard Industrial Classification (ANZSIC06), which has 94 pages devoted to manufacturing classifications, but 8 pages for financial services, 8 for healthcare, and 10 pages for professional services. This difficulty in classification extends to CPI also.

2.5.2 Technology and the free economy

An increasing proportion of our services is consumed as free or has been significantly disrupted by technology. Examples include:

- Free websites replacing the role of formerly paid products, such as online news and online encyclopaedias for information.
- Other products replaced by free alternatives, such as free GPS maps (from Google or Apple) replacing paid GPS systems
- Uber and Lyft displacing traditional taxi services, Airbnb displacing hotels
- Steam displacing traditional computer game retailers.

While most of these examples can be correctly allowed for by re-weighting indices, the process for doing this is slow relative to the pace of change in some of these industries.

Technology hardware itself is also hard to monitor for price changes. The ABS 'matches' products from month to month, so that if an existing laptop is replaced by an improved model, the discounted price of the older model is used to measure the deflationary impact of the quality improvement in the laptop. This tends to give a strong deflationary price signal, even if the amount consumers spend does not materially decrease. This effect is only corrected on re-weighting.

The quality of software is similarly difficult; Microsoft Office pricing will vary upon a new release, but it is difficult to measure how much of any price increase is due to quality.

2.5.3 Personalised pricing

More sectors have become more personalised with regard to pricing. The ABS allows for this in some areas; for example, it collects data from general insurance companies correctly to obtain premium information and collects new car prices directly from dealers. However, in other areas this is not done:

- Supermarket rewards offers
- Telco custom offers
- Negotiated discounts for some retailers
- Differential pricing (bank interest rates, professional services).

Any increase in such price differentiation makes price collection more difficult.

2.6 Main biases in CPI

There are a number of biases in the CPI. The ABS attempts corrections for some of these effects (drawn from p33 of ABS, 2011b):

- Elementary index bias: Incorrect choice and weighting of the 1,000 underlying categories used to construct CPI. For example, overweighting bread relative to beef in the index.
- Substitution bias: Ignoring substitution effects arising from a price change. For example, if a leading brand of cereal increased in price, the substitution to other cereals is not explicitly allowed for. This is partially offset by the use of geometric averaging within elementary aggregates.
- Outlet substitution bias: Ignoring shoppers who change outlets for cheaper prices (substitution to Aldi being a current relevant example).
- Quality adjustment bias: Inadequate adjustment for quality changes.
- New-goods bias: New products (e.g. smartwatches) would not be in the index at all. These can be doubly tricky as they often experience significant price deflation in their early years.
- Coverage bias: Other product categories excluded (e.g. life insurance), or geographies excluded (e.g. regional areas). Only an issue if used in a context broader than the index itself.

2.7 Future improvements to the CPI

Partly in response to the challenges above, but also recognising the opportunity offered by new datasets, there is likely to be significant change in how the ABS measures CPI.

- More regular weight updates: The ABS is considering using information from the annual Household Final Consumption Expenditure (HFCE) survey to update weights more regularly than the HES. A switch would carry its own challenges.
- More use of transaction data: Transaction (scanner) data is now used in categories representing about 25% of the CPI index. Research is ongoing on how to extend its use further – for example, whether to extend some of this beyond capital cities.

There are additional opportunities to use technology to better track inflation. In the USA, a number of organisations such as Adobe[1] and MIT[2], report their own inflation indexes, primarily using online prices and transactions. The MIT index is notable for the relatively close tracking to official CPI figures.

[1] http://www.slideshare.net/adobe/report-adobe-digital-economy-project-may-2016
[2] http://bpp.mit.edu/usa/

3 IMPLICATIONS OF THE CPI FOR USERS

3.1 Choice of index

Although the focus of the preceding chapter is on the CPI, there are a large number of indices regularly published by the ABS, which may be more appropriate gauges of inflation under different circumstances.

Main alternatives include:

- Producer price index (PPI)
- Wage price index (WPI)
- Average weekly earnings (AWE)
- Analytical Living Cost Indexes for Selected Australian Household Types (cat. no. 6463.0, see figure), which covers:
  - Employee households
  - Age pensioner households
  - Other beneficiary households
  - Self-funded retiree households

  Cost of living for age pension and beneficiary households has tended to increase a little faster than CPI over the past decade.

Figure 3 Comparison of cost-of-living indices to CPI.

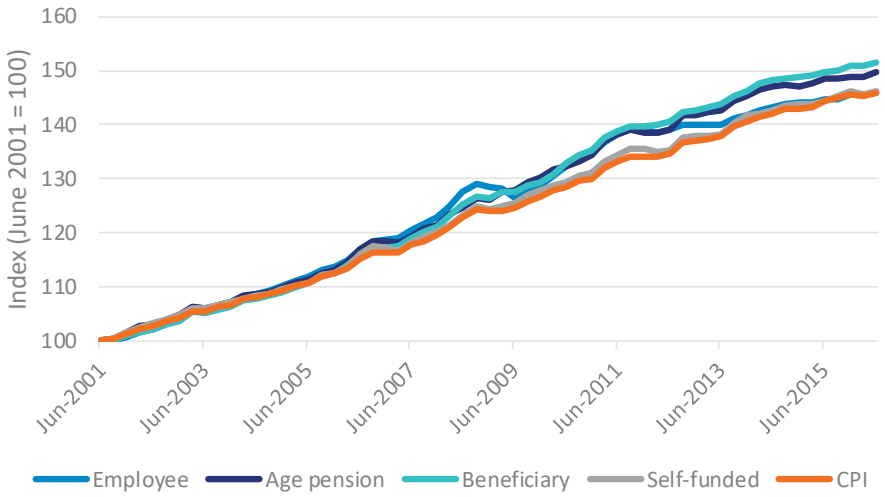

3.2 Sub-components of the CPI

Even within CPI, there are big differences between components. The chart below shows some relevant to insurance lines.

Figure 4 Price trends for health, new home price, legal services, home contents and new cars

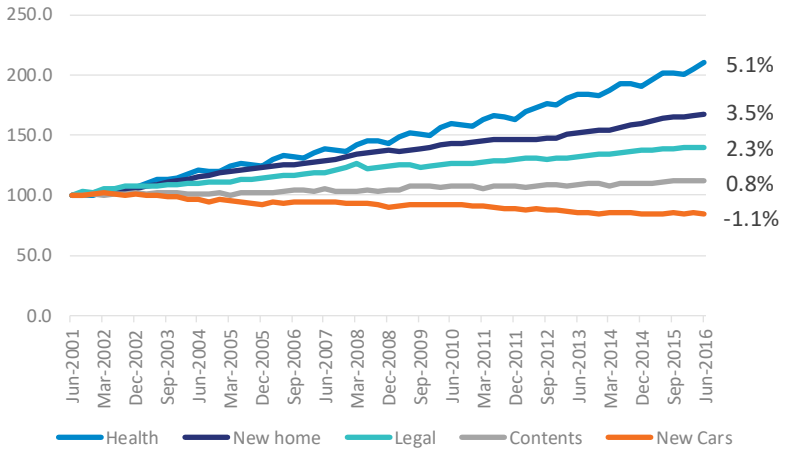

Figure 5 Price trends for subdivisions of health

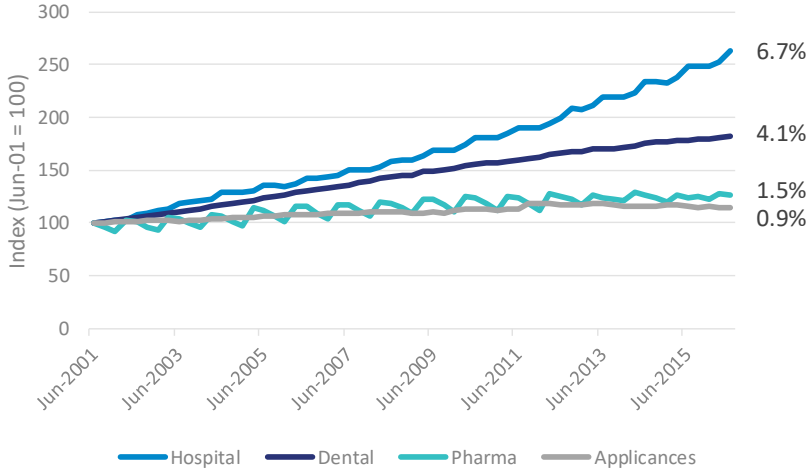

This has obvious implications. If a user has different weighting of classes for their application, the historical and future trends could be markedly different from the 2-3% range for total CPI. For example, a CPI index that excludes housing, education, tobacco and alcohol would have averaged 1.7% over the past 15 years with an implied target range of 1.2%-2.2%.

Most (if not all) insurers are aware of the differences, and allowing for different inflation for health expenses and home rebuilding costs is very common. Similarly, the exposure to currency (for imported goods) can be allowed for by correct selection of indices.

3.3 Allowance for biases

3.3.1 CPI outside capital cities

Land and rental costs have always been higher in the cities, and have tended to grow more rapidly too. This will often feed into other costs; a restaurant with lower rent may have lower prices. Regional incomes also tend to be lower too.

Figure 6 House prices for Melbourne and the rest of Victoria

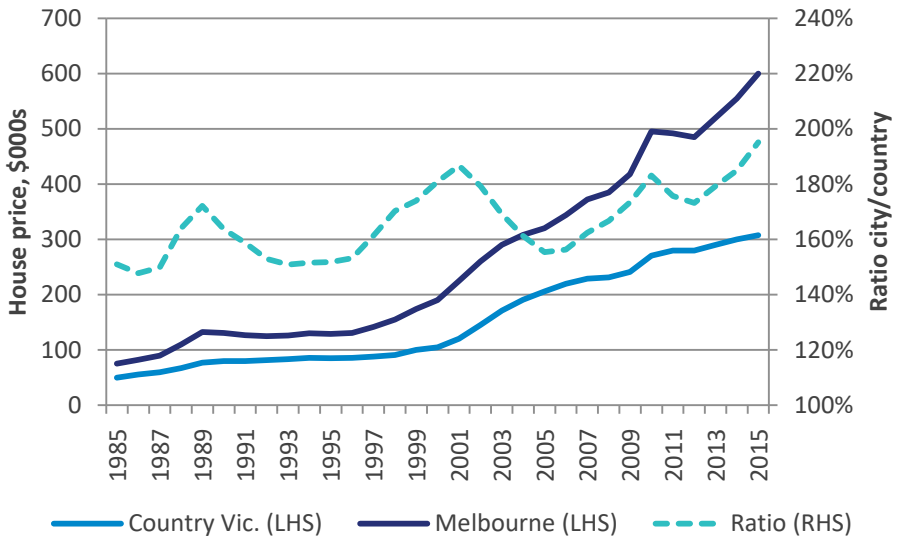

Some costs will not be lower (some may even be higher); it is worth having a think about the specific items of interest in your application.

3.3.2 Substitution bias

While there are no easy solutions estimating the substitution bias of CPI, the Reserve Bank's solution is to ignore categories that move the most in the quarter (see section 3.4). This reduces the substitution effect bias in cases of more extreme price changes.

3.3.3 Quality adjustment bias

Again, while it is difficult to easily counter the quality adjustment bias, it is possible to consider which categories typically have the most rapid quality adjustments and share of wallet increases by looking at the difference in trends between the household expenditure survey and the CPI results.

3.4 RBA target band

In setting inflation expectations, people often refer to the target bands set by the Reserve Bank of Australia (RBA). They aim for CPI to be in the 2%-3% band over the medium term. However, the RBA does not solely target raw inflation (as considered above) but considers smoothed measures of inflation – with the most prominent being the weighted median and the trimmed mean CPI. The RBA's description of the two is:

The 'Trimmed mean' is calculated by ordering all the CPI components by their price change in the quarter and taking the expenditure-weighted average of the middle 70 per cent of these price changes. The 'Weighted median' is the price change in the middle of this ordered distribution, also taking expenditure weights into account.

There are further subtleties to the implementation of these measures too. For instance, if the categories are not seasonally adjusted than some (like vegetables) have high variability and will tend to be trimmed off more often.

A natural question is whether there are systematic differences between these smoothed measures and observed CPI. This is conceivable; for example, if healthcare price increases tended to be more stable, they would more regularly feature in the central (non-trimmed) group which would bias the smoothed measures. We will focus on the trimmed mean, arguably the more popular of the two.

There has been previous research on the trimmed mean versus raw CPI. Brischetto and Richards (2006), and Roberts (2005) found that the bias of the trimmed mean was small (with various estimates, some of the order of 0.2%). The difference over the past 15 years is similarly small, with the trimmed mean 0.2% higher than the normal CPI. Over the preceding 15 years the trimmed mean was lower.

Figure 7 Quarterly CPI and trimmed inflation

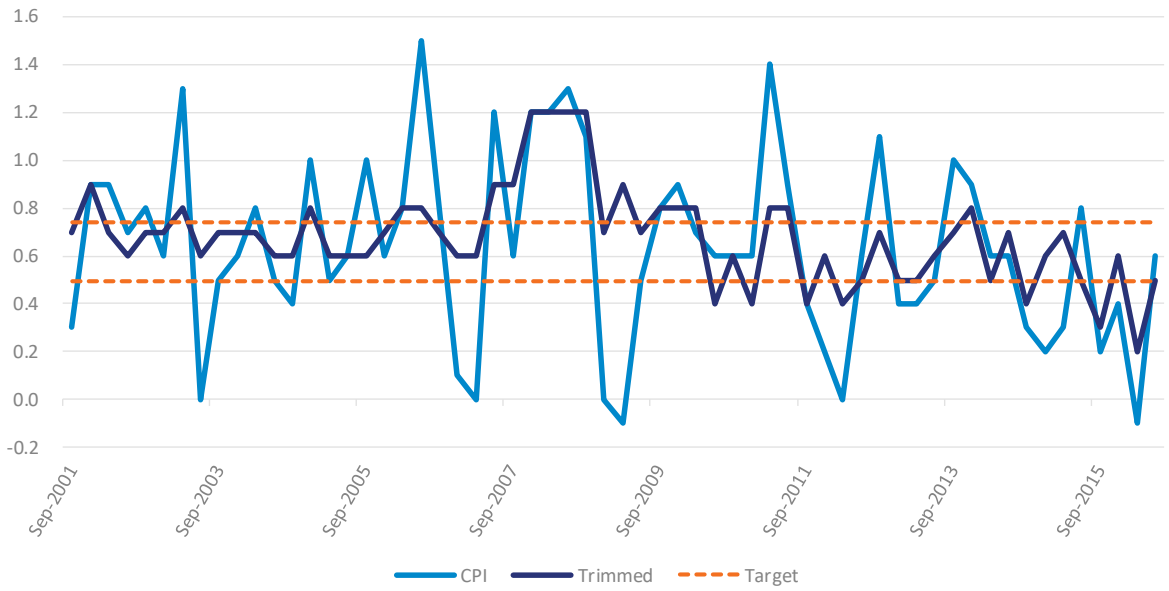

The consistency between the two means that other biases and considerations related to CPI are likely more material than the subtleties of RBA targeting.

3.5 Inflation-linked bonds

This paper does not consider inflation forecasting in great detail (see Miller, 2010) but, due to recent developments, we have reviewed the feasibility of using inflation-linked bonds. These bonds are tied to the CPI index, with coupons and face value increasing in line with index changes. In theory the difference between these bonds and nominal (non-indexed) bonds can be used to gauge market expectations on inflation. In previous decades these were thinly traded and gave only weak inflation signals as a result.

This has changed in recent years. There is now $30b (face value) of index-linked bonds issued by the Australian Government, a significant increase; it has doubled over the last five years. The longest maturity is currently 2040.

Figure 8 Face value of issued inflation-linked bonds

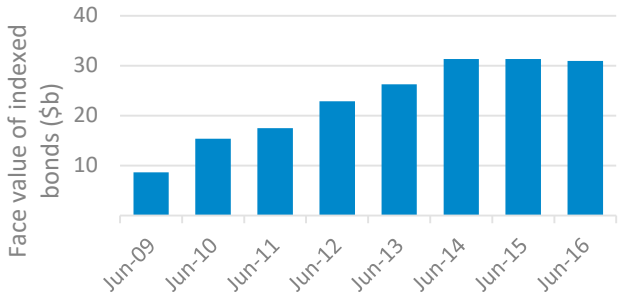

Figure 9 10-year spot yields for nominal and inflation-linked Australian Government bonds.

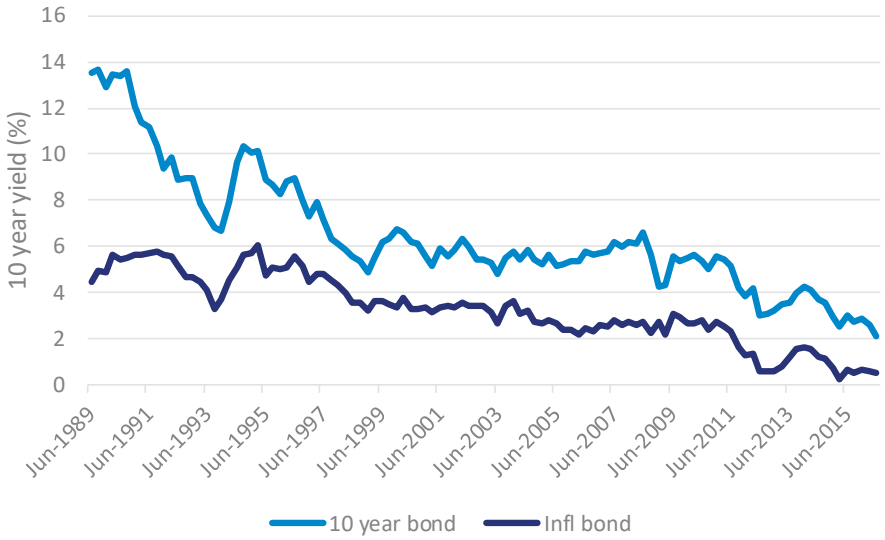

The graph shows that, for the period between June 2000 and June 2008, inflation linked bond yields remained relatively stable, while nominal bond yields drift down then up. This stability at the time suggested that inflation linked bonds were not suitable for gauging market expectations. It was unlikely that inflation expectations moved in lockstep with nominal bonds, and volumes were low – the more natural conclusion was that it was not liquid enough for use in inflation forecasting.

This appears to have changed since June 2009. The yields of the two bond types have moved together, as we would expect if yields changed more rapidly than inflation expectations. The degree of movement (the change in inflation yields is about 80% of the nominal yield change) is also consistent with other findings on the relationship between inflation and yields (see Miller, 2010). This suggests that inflation linked bonds should be reconsidered as a source of a useful market signal.

There are a couple of caveats to consider before taking the gap as a direct measure of inflation expectations:

- The total face value of issued inflation-linked bonds is still less than 10% of that of nominal bonds. This suggests that there may still be a liquidity premium attached to them.
- There is also an inflation risk premium embedded in the price of inflation-linked bonds; an additional value is attached to the benefit of removing inflation risk. It could be argued that this market pricing of inflation risk should be retained in forecasts[3], but typically insurers use a central estimate of inflation forecasts without such a premium. It is difficult to isolate this premium, as it requires another estimate of inflation expectations. Finlay (2011) attempted to measure the inflation risk premium by comparing to consensus economist forecasts. This suggested the inflation risk premium was relatively volatile, as the consensus forecasts tended to evolve gradually. Estimates of the USA and EU inflation risk premiums show variation over time, with the USA averaging close to zero and the EU closer to 0.4% over the six years to 2008 (Hördahl, 2008).

[3] The argument is analogous to that of discount rates. Observed ('market') discount rates are a combination of expected rates plus term premia. Most liabilities use the market discount rates because these better reflect the market value of future cash flows. Similarly, the difference between the market value of inflation (as measured by inflation-linked bonds) versus expected inflation requires a choice between market consistency and central estimates. The latter currently predominates, but we believe current regulatory standards would permit either.

The average implied inflation rate between June 2011 and June 2014 was 2.6%, which is close to the centre of the RBA band; this suggests the premiums discussed above have been relatively small in recent years.

4 UNDERSTANDING AWE

4.1 What does AWE measure?

Average weekly earnings (AWE) measures the take-home pay of a sample of employees for a specified week in May or November. It is done using a survey of 5,000 businesses, performed every half year. The ABS collects the number of employees (full-time and other), plus weekly earnings (total, overtime, and salary sacrifice). Bonuses are generally excluded.

AWE is an evolutionary rather than a basket based index. That is, it responds to the compositional changes in the workforce and does not attempt to compare the wages of similar jobs over time (this is the role of the wage price index, WPI). In particular, an increase in AWE could arise for a number of different reasons:

- Increase in hourly rates of pay
- Increase in number of hours worked, or overtime
- Change in the mix of full and part time work
- Change in the types of jobs being performed.

Most insurers understand these effects – often the Wage price index (WPI) is the more appropriate choice if a cost requires a similar level of expertise over time. The WPI is a 'basket of goods' approach to wages which attempts to control for changes in the nature, quality, quantity and location of work over time.

Results are split by state, gender, sector (public/private) and top-level ANZSIC industry:

B. Mining
C. Manufacturing
D. Electricity, Gas, Water and Waste Services
E. Construction
F. Wholesale Trade
G. Retail Trade
H. Accommodation and Food Services
I. Transport, Postal and Warehousing
J. Information Media and Telecommunications
K. Financial and Insurance Services
L. Rental, Hiring and Real Estate Services
M. Professional, Scientific and Technical Services
N. Administrative and Support Services
O. Public Administration and Safety
P. Education and Training
Q. Health Care and Social Assistance
R. Arts and Recreation Services
S. Other Services.

The survey was conducted quarterly up to May 2012, when it was switched to twice-yearly. It is based on a mail-out to selected businesses and refers to the week ending the third Friday in May or November. The sample is based on the ABS's list of employing businesses. It is stratified by state, sector, industry and employer size. 16% of the selected sample is rotated off each survey.

Outliers are down-weighted using the ‘surprise outlier’ technique.

4.2 What is not measured?

There are a number of groups excluded from the AWE sample, with the most notable being

- Self-employed individuals
- Proprietors of unincorporated businesses
- Defence personnel
- Those paid entirely by commission
- Those employed in agriculture, forestry and fishing (ANZSIC ‘A’).

The first two categories on the list are probably the most significant exclusions. There are about 1.25 million non-employing business, and another 0.75 million employing less than 20 people, suggesting that perhaps 2 million people fall into this category (or about one sixth of the labour force). Their omission causes some bias in estimates; households with a self-employed main income have lower than average wages (although with much greater variation than full time earnings). See Greenville et al (2013) for further detail and discussion.

5 IMPLICATIONS OF AWE INDICES FOR USERS

5.1 Why don't we use the median wage?

Income is heavily right-skewed in distribution; there are a small number of people with very large incomes that materially affect the mean. Median prices are a very common central measure for right-skewed distributions, better reflecting where 'most' observations sit. House prices, for instance, are typically reported using the median. This raises the question why we don't use median incomes in preference to, or jointly with, average income?

The main reason why means are used in the AWE survey is that it is at a business level (so the ABS collects total wages for the business, rather than a list of wages by employee). This data can be used to construct an average, but cannot be easily converted to a median figure.

The ABS does collect individual level income data, but not in a way that is as regular or robust as AWE. The household income survey measures *household* level income every two years. Another survey, the employee earnings survey was performed annually as an addition to the household labour force survey. It reported median and mean earnings for **employees**, but was discontinued in 2013.

Figure 10 Median and mean weekly earnings, employee earnings survey

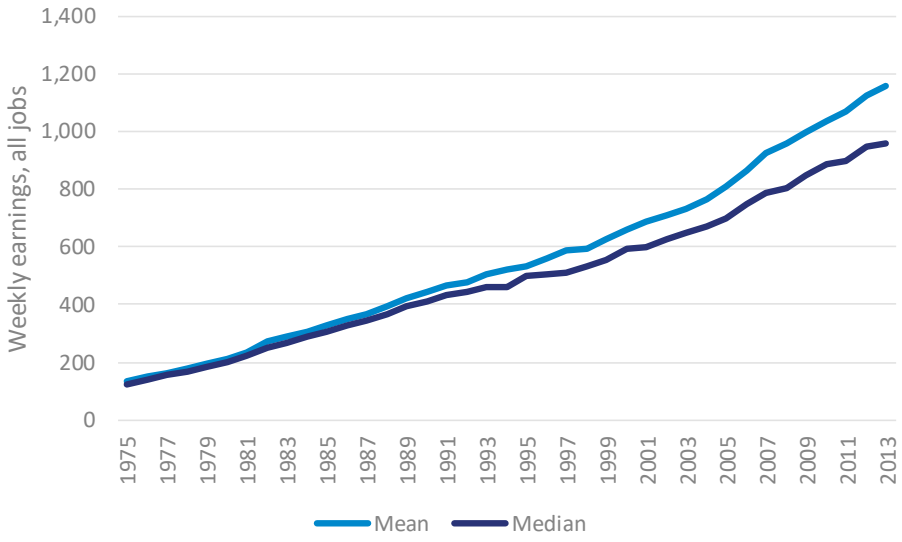

The chart shows a widening gap between mean and median incomes over time. The mean was 6% higher than the median in 1975, but was 20% higher in 2013. Virtually all this widening occurs after 1995 – the measures track each other very closely in the first 20 years of the time series. This suggests that items pegged to AWE growth for the past 20 years are about 14% higher than they would have been if median wages were used instead.

The gap between median and mean wages is one proxy for inequality. It has implications for insurers and society more broadly. For instance, it means that both workers compensation entitlements and pensions (both often tied to AWE) have increased faster than if they followed median wages.

5.2 Income inequality in Australia

5.2.1 Implications of income inequality

The distinction between median wages and mean wages only scratches the surface in terms of understanding the distribution of wages. Any widening gap between the two is indicative of more systematic growth in income inequality, a topic of particular interest to economists in the last five years.

Income inequality has a number of implications for financial services and insurance:

- As introduced in 5.1, above average growth in the top percentiles mean that 'average' wages become increasingly different from the median or even a trimmed average wage. This potentially biases AWE too high for some applications, perhaps including pensions and compensation.
- Outcomes and behaviours (such as return to work) vary significantly by income bands, so diverging income may cause composition changes in claim behaviour.
- From a customer marketing view, diverging incomes is consistent with offering the market products at different price points, perhaps under different brands.

The AWE, being an average, ignores the distribution around the average. A changing distribution can have consequences for governments and other organisations.

Internationally, increased inequality has been of considerable academic interest for other reasons (it has been equated with slower growth and increased social dislocation). See Dabla-Norris et al. (2015) for a recent international perspective from the IMF.

5.2.2 Findings on inequality income

Income distributions are more difficult to measure, and so their availability in Australia is somewhat limited. The best sources of income distributions are the:

- Household Expenditure Survey (HES), which is only completed every five or six years and most recently in 2010.
- ABS Survey of Income and Housing (SIH), which is completed every one or two years and reconciled to the HES.

The time series are intermittent and not particularly up-to-date.

In Australia, there is some mixed evidence of increased income inequality. The figure below shows the income and wealth growth in Australia over the past decade or so. The income shown is 'equivalised' real disposable income, which the ABS has developed to allow for tax, inflation, cost of living and household size. At a quintile level, income gains have been reasonably even, with only the top quintile having any visible difference. The top quintile has seen income growth of 80% over 19 years, compared to 60% for other groups. The Gini index (a common measure of income inequality) grew only slightly, from 0.30 to 0.33 over that period.

The story is very different for household wealth (right panel), with the top quintile seeing far more rapid growth. Much of this is tied to the housing market, Australia's primary source of household wealth; gains have benefited those with property already.

Figure 11 Equivalised disposable real income growth (left) and real wealth growth (right) for income and wealth quintiles

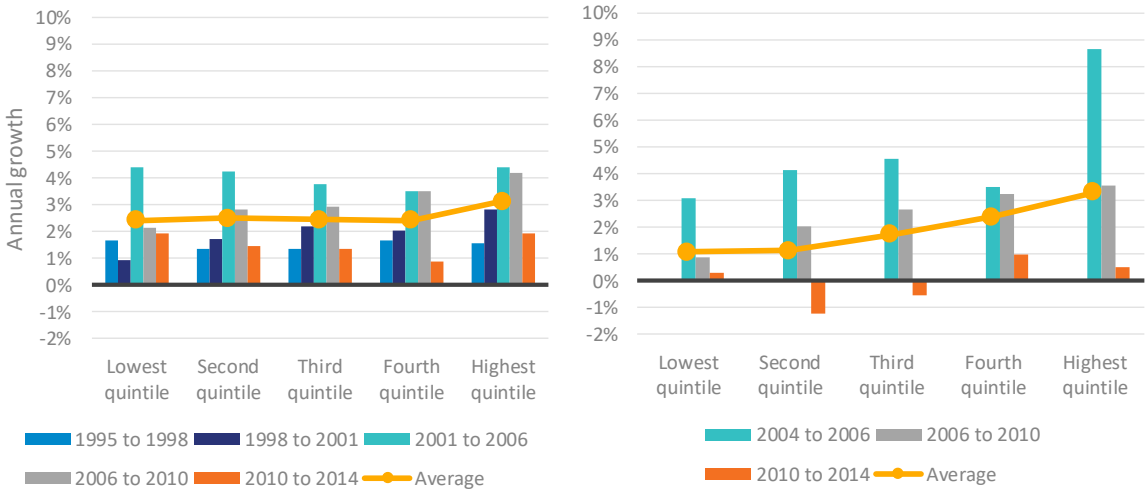

Greenville et al (2013) also give a good description of inequality in Australia, using individual level data from the HES. They find more disparity in wage growth when looking at gross individual wages by decile. This suggests that the household measures in the figure above obscure some of the divergence seen in the wage economy, particularly for fulltime employees and self-employed people. This is partly explained by reasonable income growth for people on benefits such as the age pension, but also by increased employment in lower income households (see section 5.3). Taxes are also a little less equalising that they were. They estimate that wage inequality, measured via the Gini coefficient, has increased from 0.35 in 1988-89 to 0.41 in 2009-10. They also note that the difference between gender pay amounts to about 15% of observable wage inequality.

Figure 12 Change in the distribution of real labour income, 1988-89 to 2009-10. Taken from Greenville et al (2013).

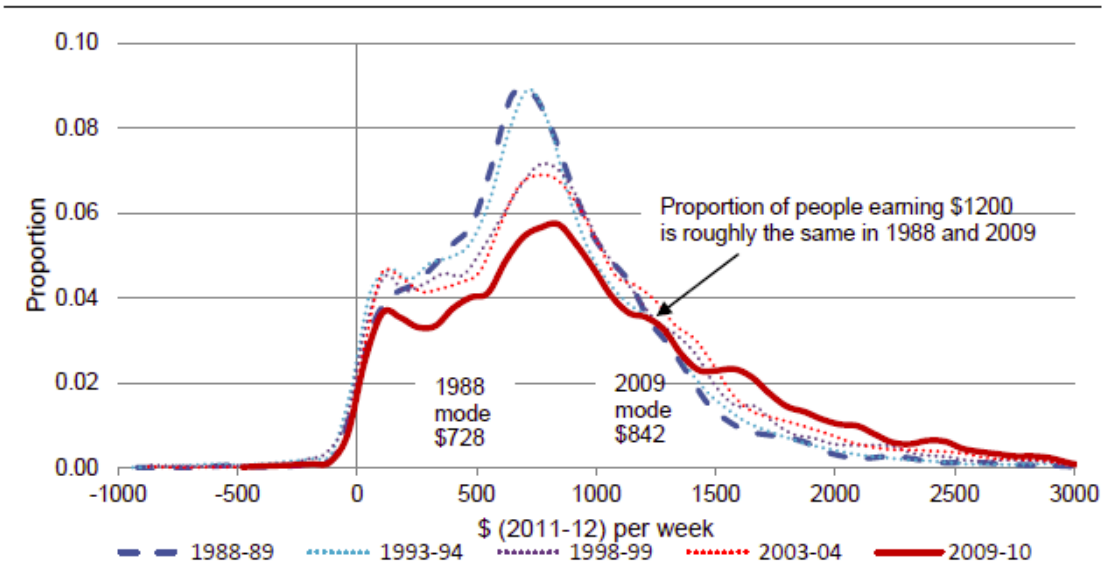

There is also other more anecdotal evidence of income inequality at the top end. For instance, CommSec maintains an index of luxury car vehicle sales. These have increased rapidly, particularly in the past few years.

Figure 13 Annual luxury car sales, CommSec, August 2016[4]

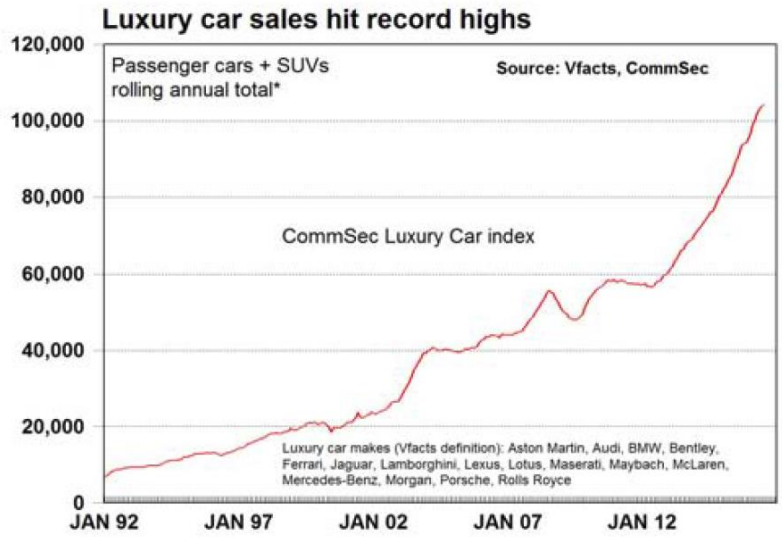

5.3 Is household wage or individual more relevant?

While individual wages are probably the most useful when considering inflation across the economy, it is worth recognising that economic activity generally operates at a household level. Household incomes have increased somewhat faster than individual wages due to compositional reasons. Over the 10 years to 2009-10, the average number of adults employed per household rose from 1.21 to 1.30. Or equivalently, the proportion of adults in employment rose from 56% to 60% (HES data). This means that household incomes have been roughly 0.7% higher per annum as more households earn a second income. Further, the effect is more pronounced in lower income households; thus the individual income divergence (where upper deciles grow faster than lower ones) is partially offset by higher employment rates in lower deciles.

AWE also ignores the impact of other factors, such as tax, which are sometimes significant; if all the income growth is for people on higher tax rates, the average after-tax growth rates will be lower.

5.4 The switch to half-yearly AWE

In May 2012 the ABS switched from quarterly AWE updates to half-yearly as a cost-saving measure. This section considers the implications of this on some actuarial work. The main consequence is that AWE can be up to nine months out of date when applied (including the reporting lag), so a later update to can cause larger revisions when applied.

Consider a simplified reserving example. Suppose a reserving exercise does the following:

- Claims are inflated to June 2016 values using the February 2016 release (November 2015 salaries). An assumed level of inflation between November and May is set using a historical average.
- A reserve is set, using averages of inflated claim levels.
- Claims are reinflated in August 2016, after the release of the May 2016 salaries. The 'change' in levels is the difference between average and actual change over the six months to May 2016.

[4] https://www.commsec.com.au/content/dam/EN/ResearchNews/ECO_Insights160816-record-luxury-car-sales.pdf

» No other assumption changes are made to the reserve.

In this case the reserve estimate changes by the difference between actual and expected AWE inflation over six months, since this flows directly into claim inflation prior to setting other assumptions. So if 2.0% inflation was expected, and 3.0% over six months observed, then a 1.0% increase in liability is observed.

Previously this experiment would have depended on a three-month gap, so we can estimate the additional volatility introduced by the change in sampling frequency. We can use the quarterly series before May 2012 to compare how a 3-month lag and 6-month lag affects the volatility. We call the difference between the actual AWE change and the historical average the AWE 'surprise'. We use the NSW seasonally adjusted index, full time ordinary time earnings.

The results are shown in the figure below. The increase in AWE surprise is evident; the standard deviation grows from 1.0% to 1.4%, and the width of the 5th to 95th percentile range grows by half, from (-1.5% ,1.6%) to (-2.2% to 2.4%). The fact that 2% surprises are now not uncommon means that there are significant volatility implications from using AWE in a manner similar to this reserving example.

Figure 14 Quarterly versus half-yearly AWE 'surprise' relative to historical average. NSW seasonally adjusted series

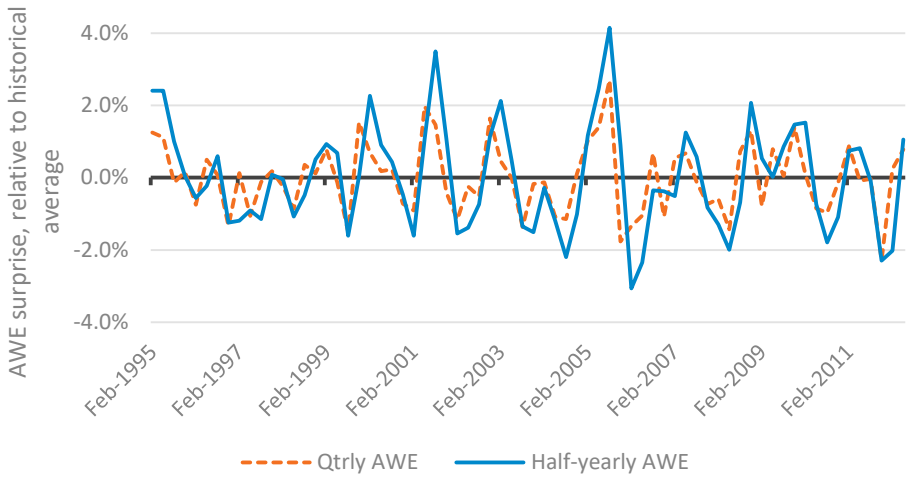

We note:

» The forecast used here is the historical average of AWE. If good six-month forecasts of AWE existed, then this would reduce the average magnitude of surprises. Our anecdotal experience with various forecasts is that the benefits of this mitigation are limited.

» There are ways to restructure this experiment (and similar actuarial work) to reduce the volatility. For instance, a modified AWE series can be used which is more heavily smoothed, reducing the impact of the latest point estimate.

» As an extreme example of the previous bullet, we could replace the AWE index inflation with a fixed annual assumption (of 4% say), which removes the issue entirely.

» Another obvious alternative is to switch to CPI with extra superimposed inflation. While AWE and CPI can have different drivers, the correlation is strong enough that the change is plausible. CPI is still reported quarterly, and tends to be less volatile than AWE anyway.

To test the second bullet, we created a modified AWE index (for NSW seasonally adjusted) at year t:

$$AWE_{mod}(t) = \frac{1}{2}AWE(t) + \frac{1}{2}(1+i)^{0.5}AWE(t-0.5)$$

Here i is a long run AWE inflation rate (4.3% for this example). We then change the reserving setup so that claims are inflated using the modified index. This introduces some bias (if wages have genuinely leapt in a quarter, then the modified index will understate the impact), but reduces noise (the half-yearly AWE series in this example has a correlation of -0.28 with the previous change). As you would expect, using this modified index significantly reduces the surprise in a valuation. This is shown in the figure below. In fact, the standard deviation of the surprise is about 10% lower than the quarterly series. The width of the 5th to 95th percentile range shrinks to (-1.4%, 1.5%), again smaller than the original quarterly surprise series. This reduction is significantly larger than that obtained by switching to the trend estimate (which is also subject to historical revision).

Figure 15 Quarterly and half-yearly AWE 'surprise' relative to historical average, plus the surprise if the modified AWE series is used instead. NSW seasonally adjusted series

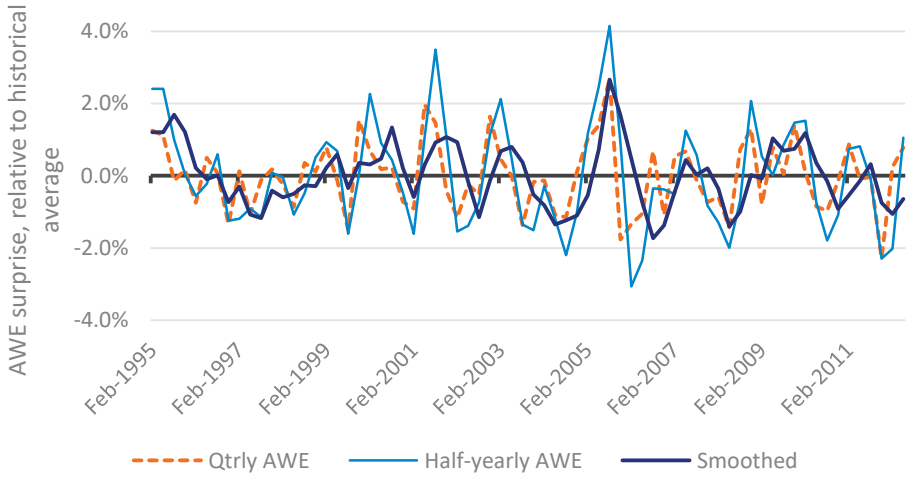

This suggests that heavier smoothing is feasible as a way of mitigating the switch to half-yearly AWE.

6 CONCLUSIONS

So where does this leave us? Perhaps not much better off in terms of actuarial best practice, but the investigation raises a number of points that are worth repeating:

- Inflation indices are highly definitional and reasonably complex to construct. These mean that they are not always the most appropriate index to apply, and perhaps we should have more confidence in altering the index to suit a particular need.
- Data collected and reported by the ABS is rich, despite budget pressures and limitations on collection frequency. A knowledge of what is available helps broaden understanding. For example, the subcomponents reported along with CPI are very useful.
- Technology and innovation pose both a challenge and opportunity to robust data collection. Inflation indices are slow to reflect the digital economy and its role in prices.
- Inflation-linked bonds are more useful as a market estimate of future inflation now than they were five years ago, thanks to a marked increase in issuance.
- By focusing on the average, AWE inflation is sensitive to inequalities in wage growth. The median wage historically has grown more slowly, and in some contexts may be a more appropriate index of change.
- There are some significant considerations when moving between individual gross wages and household income. The differences may be important for some financial service contexts.
- The switch to half-yearly AWE reporting by the ABS causes difficulties for some types of actuarial work. Work-arounds should be considered, including the use of more heavily smoothed AWE indices.

7 REFERENCES

Note: All weblinks accessed in September or October 2016 unless otherwise stated.

K. C. Ahlgrim, K. C. and D'arcy, S. P., 2012. The Effect of Deflation or High Inflation on the Insurance Industry. Casualty Actuarial Society 2012. https://www.soa.org/Files/Research/Projects/research-2012-02-effect-deflation-report.pdf

Australian Bureau of Statistics, 2011a. A Guide to the Consumer Price Index: 16th Series. http://www.abs.gov.au/ausstats/abs@.nsf/mf/6440.0

Australian Bureau of Statistics, 2011b. Information Paper - Consumer Price Index: Concepts, Sources and Methods. http://www.abs.gov.au/ausstats/abs@.nsf/mf/6461.0

Australian Bureau of Statistics, 2011c. Household Expenditure Survey, Australia: Summary of Results, 2009-10 http://www.abs.gov.au/ausstats/abs@.nsf/mf/6530.0/

Australian Bureau of Statistics, 2014. Average weekly earnings and wage price index – what do they measure? http://www.abs.gov.au/AUSSTATS/abs@.nsf/Previousproducts/6302.0Main%20Features9May%202014

Australian Bureau of Statistics, 2015a. Survey of Income and Housing, User Guide, Australia, 2013-14. http://www.abs.gov.au/ausstats/abs@.nsf/mf/6553.0

Australian Bureau of Statistics, 2015b. Household Income and Wealth, Australia, 2013-14. http://www.abs.gov.au/ausstats/abs@.nsf/mf/6523.0

Australian Bureau of Statistics, 2015c. Information Paper Introduction of the 16th Series Australian Consumer Price Index. http://www.abs.gov.au/ausstats/abs@.nsf/mf/6470.0

Australian Bureau of Statistics, 2015d. A guide to understanding average weekly earnings statistics. http://www.abs.gov.au/ausstats/abs@.nsf/Lookup/6302.0main+features8Nov%202014

Australian Bureau of Statistics, 2016a. Enhancing the Australian CPI: A roadmap. http://www.abs.gov.au/ausstats/abs@.nsf/mf/6401.0.60.001

Australian Bureau of Statistics, 2016b. Increasing the Frequency of Consumer Price Index Expenditure Class Weight Updates. http://www.abs.gov.au/ausstats/abs@.nsf/mf/6401.0.60.002

Brischetto, A., and Richards, A,2006. The performance of trimmed mean measures of underlying inflation. RBA research discussion paper, http://www.rba.gov.au/publications/rdp/2006/pdf/rdp2006-10.pdf

Dabla-Norris, M. E., Kochhar, M. K., Suphaphiphat, M. N., Ricka, M. F., & Tsounta, E., 2015. Causes and consequences of income inequality: a global perspective. International Monetary Fund. https://www.imf.org/external/pubs/ft/sdn/2015/sdn1513.pdf

De Ravin, J. and Fowlds, M., 2010. Inflation risk in general insurance, General insurance seminar 2010. http://www.actuaries.asn.au/Library/Events/GIS2010/GIS10_Paper_Deravin.pdf

Finaly, R., 2011. Estimating inflation expectations with a limited number of inflation-indexed bonds. RBA research paper, http://www.rba.gov.au/publications/rdp/2011/pdf/rdp2011-01.pdf

Greenville J., Pobke, C. and Rogers, N., 2013. Trends in the Distribution of Income in Australia. Productivity Commission Staff Working Paper. http://www.pc.gov.au/research/supporting/income-distribution-trends/income-distribution-trends.pdf

Hördahl, P., 2008. The inflation risk premium in the term structure of interest rates. BIS Quarterly Review, September 2008. http://www.bis.org/publ/qtrpdf/r_qt0809e.pdf

Miller, H, 2010. Towards a better inflation forecast. General insurance seminar 2010. http://www.actuaries.asn.au/Library/Events/GIS/2010/10.c%201015%20Miller.pdf

Pearson, E. and Beynon, C., 2007. Superimposed inflation – Australian Accident Compensation Landscape Accident Compensation Seminar 2007. http://actuaries.asn.au/Library/p.2_ACS07_paper_Pearson_Superimposed%20inflation_paper.pdf

Roberts, I. (2005). Underlying inflation: concepts, measurement and performance. RBA research discussion paper, http://www.rba.gov.au/publications/rdp/2005/pdf/rdp2005-05.pdf