Towards a big data CPI for New Zealand

Paper presented at the Ottawa Group 2017

Eltville, Germany

Alan Bentley and Frances Krsinich

Abstract

In our digital age, what's the best way to measure inflation? There is an abundance of new data sources. Some of these will prove useful, given the right mind-set and methods, others may prove distracting or prohibitively hard to obtain access to. In this paper we take stock and explore the path to a 'big data' consumers price index (CPI) for New Zealand.

We discuss the benefits of a big data approach, reflecting on our early adoption of model-based approaches to price measurement, such as using a hedonic model for second-hand cars, and retail transaction – scanner data – for consumer electronics products (including TVs, computers, and digital cameras) in the CPI. A big focus of Stats NZ's work in this area has been the development and theoretical justification of a regression-based approach called 'the FEWS index' that will produce non-revisable quality-adjusted indexes even when there is no explicit information on product characteristics.

We look towards the future by considering the opportunities that are currently in front of us. Notably, Stats NZ recently signed up to the purchase of a year's-worth of daily web-scraped online price data from PriceStats, the commercial counterpart of Massachusetts Institute of Technology's (MIT's) Billion Prices Project. This data captures, in real-time, online prices for a wide range of different New Zealand retailers and it will enable us to do detailed research on the potential for enhancing and improving our current data collections and price measurement. Using the online data in combination with expenditure information from surveys or scanner data presents a rich opportunity for more frequent and timely price indicators than are currently available.

Preliminary research on measuring rent price change from administrative data is used to illustrate our approach to the art of using big data. This data source highlights the opportunities and obstacles of coverage, timing, and quality adjustment.

Developing a big data mind-set

Ideal data

The ideal dataset to measure CPI inflation, would have real-time data covering price and volume (quality and quantity) information for all products (goods and services) acquired by private New Zealand resident households whilst in New Zealand (online and offline outlets).[1] Figure 1 shows a visualisation of this ideal dataset, using three dimensions to define the target population of transactions – products, by place, by time.

Figure 1

Target population of transactions
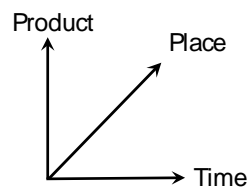
No comprehensive dataset like this exists. We have an estimation problem: to measure inflation with limited data.

Designed data and price index methods

Historically, data sources for CPI measurement are commissioned sample surveys. The underlying philosophy of price measurement, and many of the practicalities and conventions, adopted at the start of the New Zealand CPI, over 100-years ago, still form the bedrock of price measurement today. For example, the New Zealand convention of using index numbers with a reference of 1000 can be traced back to 1914. Appendix A elaborates on the similarities between contemporary methods, and those adopted in 1914.

Tracking price change for a fixed basket of commodities, using a Laspeyres-type index formula, has been the foundational paradigm throughout. Yet the methods of data collection and price index methods, to control for quality change, were designed for the customer needs and economy of the early 20th century.

Over the past century, there have been many successful advances to CPI data capture and measurement. The basket of products tracked has expanded to about 700 goods and services, representing expenditure of all kinds (target coverage of products progressively expanded from "ordinary necessaries" in 1914 to reach near universal coverage by 1975 (Bentley, 2014)). Geographical coverage reached a height of collection in 25 urban areas in 1965 (Bentley, 2013). This has been scaled back in recent years to 12 regional centres, reflecting a greater homogeneity of prices (Statistics NZ, 2014). Quality adjustment methods have become

---

[1] This paper will refer to the New Zealand consumers price index (CPI) for illustrative purposes. The same underlying dataset is used for the household living-costs price indexes (HLPIs), which show the experience of inflation for different groups in society (see Statistics NZ, 2016a). The discussion of new methods and data sources can easily be extended to other price indexes, such as those covering businesses and overseas trade.

increasingly important as products have become more complex and product life-cycles have decreased. Krsinich (2014a) gives a history of quality adjustment methods in the NZ CPI.

Retail transaction data, commonly referred to as 'scanner data', has been used to supplement Household Economic Survey expenditure data, for calculating CPI expenditure weights, since 2006 (Statistics NZ, 2006). Stats NZ was an early adopter of hedonic methods with our introduction of a hedonic price index for used cars in 2001 (Krsinich, 2011). At the end of 2014, New Zealand became the first country in the world to incorporate direct measurement from transaction (scanner) data for consumer electronics products into our CPI (Krsinich, 2014b). Yet, all of these advances have been supplementary and evolutionary. We have not yet made the leap to fully reap the benefits of all of the data available today. Nearly 20,000 prices were collected each quarter in 1914 compared with about 100,000 today (excluding retail transaction data). Whilst sample size is not the only determinant of quality, a 5-fold increase in data used seems small in comparison to changes in product sophistication and computer processing power to handle larger datasets over the same timespan.

Found data and model-based methods

Today, an abundance of new data are created in real-time as a digital footprint of our day-to-day lives. We argue that to take best advantage of these advances requires a paradigm shift away from a primary focus on sample surveys and index number theory to a 'big data mind-set'. Big data (transaction, online and administrative data) is 'found data' in the sense that measuring CPI inflation is a secondary use of the data – the data was not created with this use in mind. Coverage and access become key issues. Automated, scalable, price measurement methods are a must.

Figure 2 shows a stylised view of how big data differs in coverage of the CPI target population compared with traditional data, such as designed sample surveys.
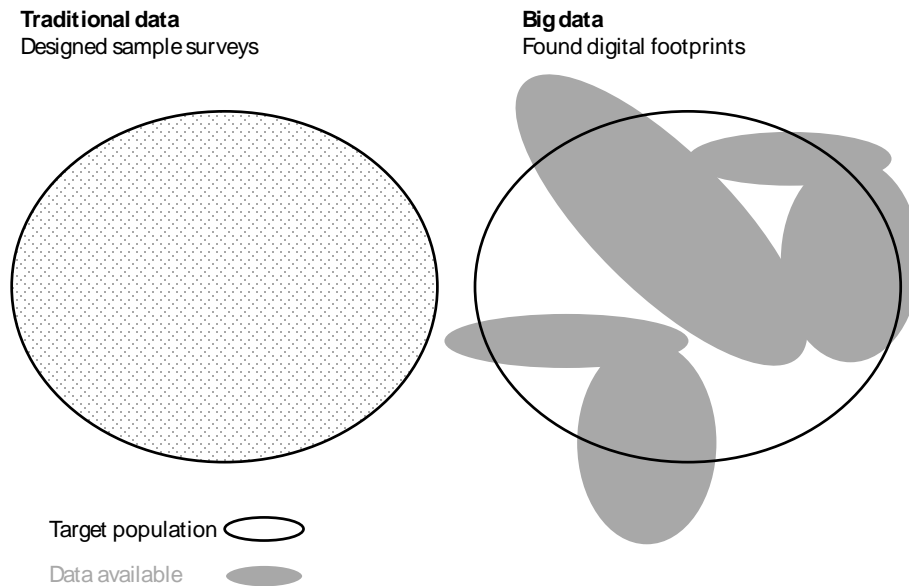
Traditional data is typically small sample surveys, well aligned with the target population.[2] Coverage of products is achieved by selecting goods and services to include in the CPI basket that are well spread across the expenditure classification. Sample design criteria include consideration of the place of transaction: geospatial characteristics such as region, and outlet type (supermarket, department store, appliance retailer etc.). Collection is expensive, so generally temporal coverage is infrequent. In New Zealand, current collection is at most weekly, and generally monthly or quarterly. See Bentley (2013) for a more comprehensive description of current data collection methods.

Big data tends to have very comprehensive coverage for a somewhat haphazard subset of transactions. It's easy to be wowed by the size of the data, such as the number of transactions, or the velocity of daily price observations. Yet, it is also important to consider coverage: under-coverage (for example missing products, outlet types, service providers) and over-coverage (business or tourist expenditure) of the found data, and make a judgement about the significance of this.

---

[2] Many CPI collections are not scientifically random designs (random probability of selection), yet they are designed in such a way as to be close proxies to such an approach (see ILO et al, 2004).

Figure 2

Paradigm shift: traditional to big data

Traditional data
Designed sample surveys

Big data
Found digital footprints

Target population

Data available

Drivers for change

Competing measures of price change

A reality of our digital age is that we are no longer the sole provider of information about price change. In recent years we have seen an expanding list of alternative measures of price change, for example:

- Rent prices from TradeMe's rent price index (TradeMe, 2017) and Ministry of Business, Innovation and Employment's tenancy bond administrative data (MBIE, 2017)

- Petrol prices from NZ Automobile Association's Petrolwatch (AA, 2017) and CardSmart's Pricewatch (CardSmart, 2017)

- ANZ monthly inflation indicator (ANZ, 2017)

- Billion Prices Project daily inflation measures (Cavallo, A & Rigobon, R, 2016).

- Adobe Digital Price Index (Adobe, 2017)

A key dimension to all these data sources is that they are more frequent than our quarterly CPI. The latter two data sources do not yet cover New Zealand, but there is no reason to believe similar data could not become available in the near future.

In our role as New Zealand's national statistics office (NSO) we are entrusted by the government and people of New Zealand to provide reliable official statistics that are the result of a careful statistical production process (Statistics NZ, 2016b). Yet we can no longer assume, if we ever could, that simply because our inflation statistics are produced under our brand and remit as NSO that they will be considered the best measure of inflation.

Adopt, adapt or explain

As alternative, digital-footprint, data and statistics about prices become available we need to be able to rapidly evaluate their merits or pitfalls. We need to be able to adopt, adapt or explain. For inflation measurement a critical aspect, in addition to coverage, is quality adjustment. That is, to control for changing product characteristics over time to capture pure price change. Whether or not alternative data sources simply reflect changes in average prices, or are price indexes that control for quality change, is often unquestioned. Understanding and explaining differences in coverage and methods is a key starting point. This may be sufficient. If alternative measures are well understood, complementary data about prices (or inflation) then this is to be welcomed. The focus of this paper is to consider whether it would be beneficial to adopt or adapt alternative data for official inflation measurement. This may include short-term indicators of inflation that are timelier reads for a subset of transactions.

Monthly CPI

There is a known desire for timelier, higher-frequency, reads of inflation from some CPI customers. The need is perhaps more acute in New Zealand given the CPI is produced on a quarterly frequency. The 2013 CPI Advisory Committee recommended that the CPI be produced monthly to benefit monetary policy setting and align with international best practice (Statistics NZ, 2013). Data collection costs have, so far, precluded a scaled-up, survey-based, approach to increased frequency. Public consultation, following the 2013 committee, found a more mixed view across the wider community on whether the benefits would outweigh the increased costs of production, using existing data collection methods (Statistics NZ, 2014).

Efficient public services

A low-cost approach to increased frequency, using big data, is a smart approach that should wash well with CPI customers and the wider community. New Zealand's Better Public Services initiative is challenging all public service organisations, including Stats NZ, to "change, develop new business models, work more closely with others and harness new technologies in order to meet emerging challenges" (State Services Commission, 2016). Result 10 – New Zealanders can complete their transactions with government easily in a digital environment – provides an aspiration for digital-first measures of inflation.

In our digital-age many expect Stats NZ's data capture methods to include digital transactions, such as online purchases, in order to be reflective of the modern world and to reflect the Better Public Services aspirations for a modern and efficient public service. Physically visiting retail outlets and undertaking postal surveys are expensive activities which may look old fashioned. We need to continually challenge whether the costs, including respondent burden, can be justified. Where this is the best mode of collection we need to be able to clearly articulate the reasons why digital-data can't be adopted.

International best practice

Aligning with international best practice has long been of importance to the New Zealand CPI. The foundational methods adopted in 1914 were greatly influenced by those in Australia. The rationale for consistent approaches, to aid comparability and economic reasoning, are as relevant as ever.

Historically, in New Zealand and other countries, index number methods start with a simplifying assumption of a static universe of products over time (ILO et al, 2004). Subsequently, methods have been developed to appropriately quality-adjust when products do change over time. Bentley (2011) and Krsinich (2014a) describe the variety of approaches currently used in New Zealand. These depend on data source and product sophistication. This approach of making ad hoc 'quality adjustments' is often time consuming, can be subjective, and is not entirely satisfactory from a consistency view point.

'Hedonic', or regression-based methods have evolved from the use of big data such as scanner data to determine quality adjustment factors when replacing products in a fixed-based context, to more automated model-based approaches where, for example (in the time-dummy hedonic method), the price index is derived directly from the parameters on time when price[3] is modelled against time and the characteristics of products.

de Haan and Krsinich (2017) show that time-dummy hedonic indexes can be understood in the context of index theory as a geometric version of a quality-adjusted unit value index.

Krsinich (2016) has shown that the FEWS index - or a time-product dummy (TPD) index with a window-splice - is equivalent to a fully-interacted time-dummy hedonic index, if all the price-determining characteristics are expressed as categorical variables. The window-splice incorporates a revision factor to ensure that the index is updated to reflect the impact of new products and the improved estimates of the products' fixed effects with an increased estimation window.

Internatonally, there is a growing recognition that 'multilateral methods'[4] are required for optimal and robust price measurement from big data (Australian Bureau of Statistics, 2016).

---

[3] Usually, log of price.

[4] Multilateral methods are those that incorporate a full estimation window of data – usually at least a year – rather than referring only to the current and base period quantities and prices, as bilateral indexes such as the Laspeyres or Tornqvist indexes do. Multilateral methods are often model based (eg time-dummy hedonic, ITRYGEKS, FEWS) however the GEKS method is an example of a multilateral method which is not model-based.

Coverage of household expenditure by big data

A good place to begin assessment of the coverage of the digital data we find is the Household Economic Survey – New Zealand's household budget survey. This designed sample survey, of about 3,000 households, provides excellent cross-sectional coverage of the target population of transactions. Conducted every third-year, the product and place dimensions of our ideal data are well covered. Being a scientifically random sample of households we can estimate total household expenditure, and, more importantly for inflation measurement, expenditure by products and place (both outlet type and region).

Product coverage

Comparing the coverage of household expenditure, as used for CPI weights, [5] with online prices collected as part of our collaboration with PriceStats is shown in Figure 3. The data covers 13 retailers in the initial trial. With just 13 retailers we obtain 160,000 daily price observations and coverage in about half of all detailed COICOP (Classification of Individual Consumption According to Purpose) categories – 49 out of 98 level 3 COICOP categories; 24 out of 43 level 2 COICOP categories (COICOP level 2 coverage is shown in the Appendix B). The websites chosen in the trial were selected to understand a range of different site designs in addition to providing good coverage of products.

Figure 3

Online data trial:
Better coverage of goods than services
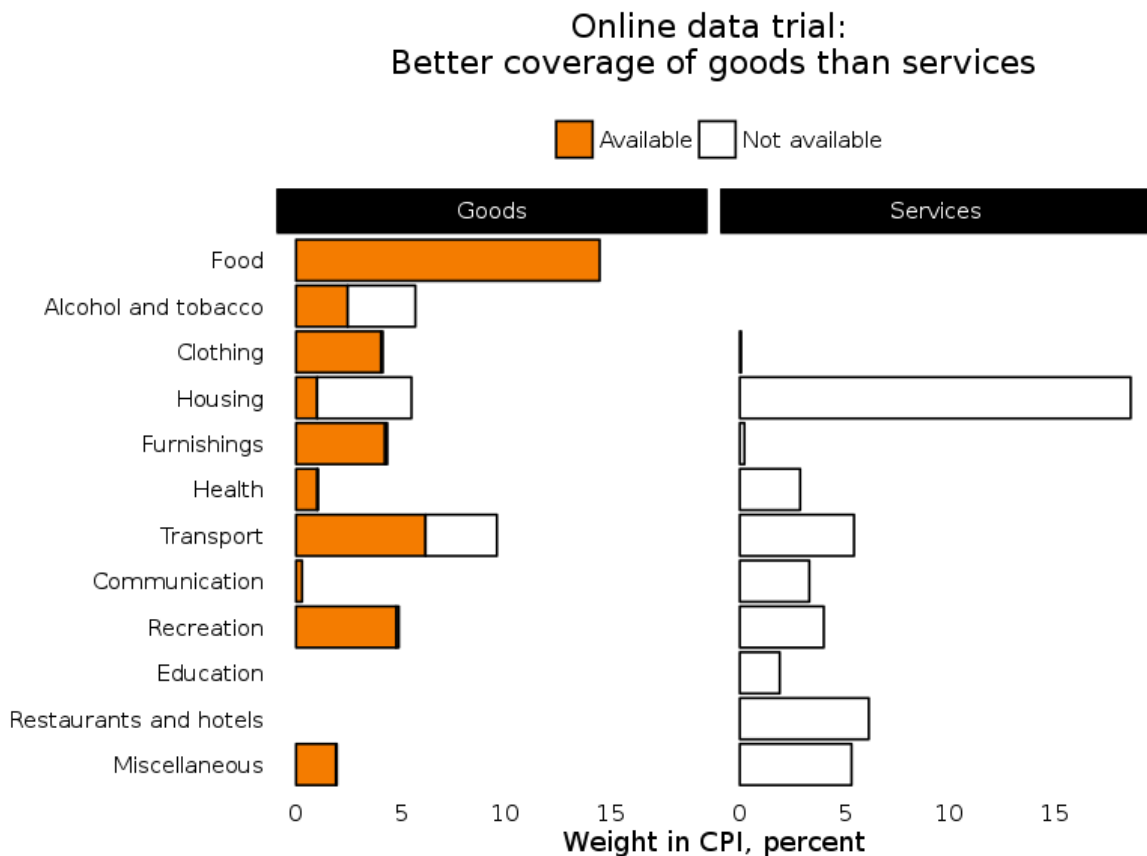
Source: Stats NZ | PriceStats

We find that coverage is much better for goods than services. This is both heartening and disappointing. It's heartening as goods are predominately what is currently priced by physically

---

[5] Household Economic Survey product-level expenditure estimates are triangulated against other available data sources as part of the creation of CPI expenditure weights (Statistics NZ, 2014).

visiting stores, which is an expensive mode of data collection used for a large number of products. It's disappointing as we will need to look to other sources to get good coverage of services. Online data alone is unlikely to provide good coverage for many services as, currently, many service prices are not available online.
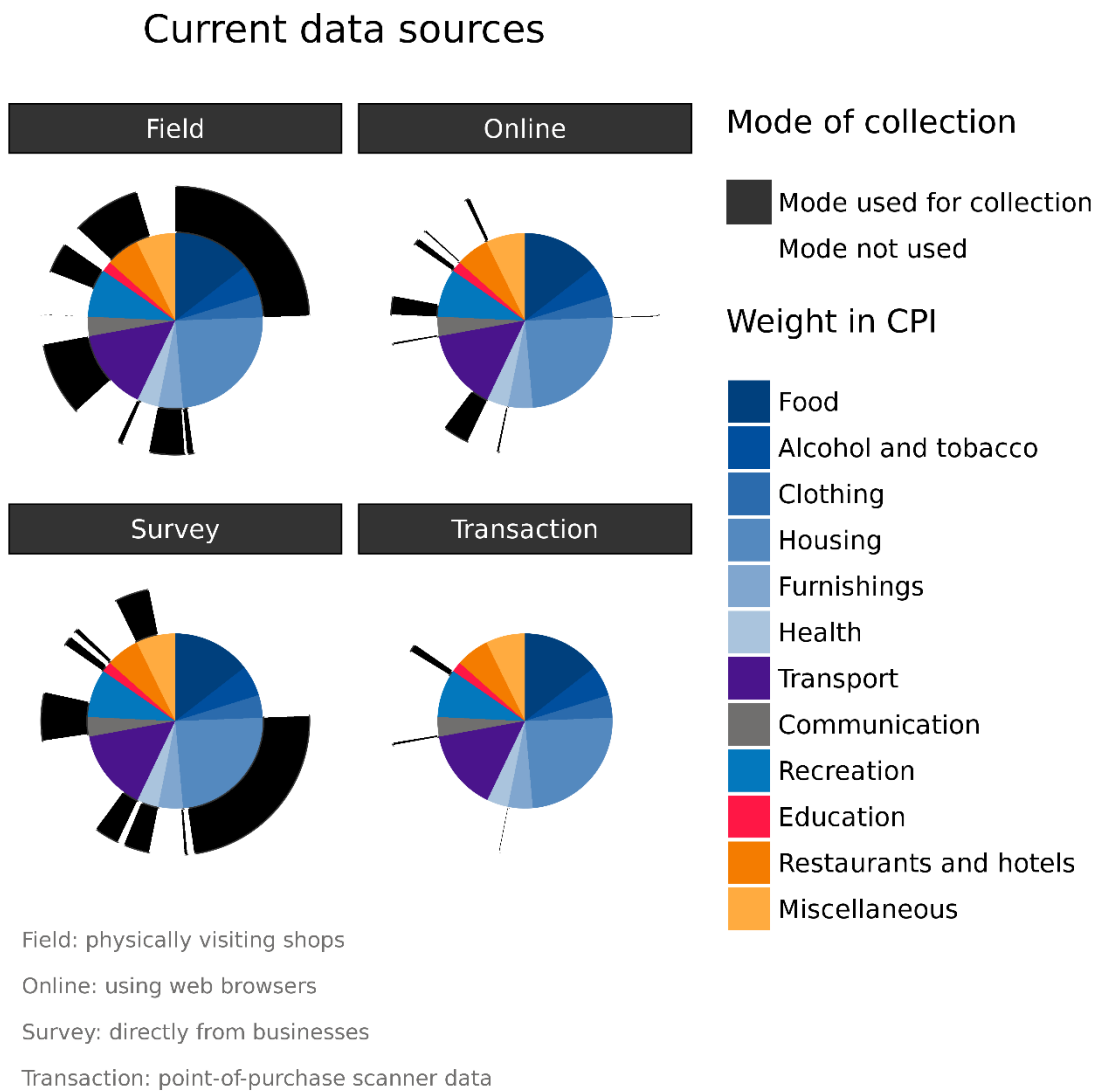
Looking at other potential big data sources, shown in Table 1, coverage of most goods is likely to be achievable. We have been using transaction data for consumer electronics products in our CPI since 2014. Access has been obtained by using a third-party company who also perform some data cleaning and aggregation. Our innovative approach for these products is explained in Case Study 1.

Table 1

Actual and potential big data sources

Source	Access	Product coverage	Observations
Transaction data	Third party	Consumer electronic products	Used in CPI since 2014 – see Case Study 1.
Online	Currently collected online	Various, including: Airfares, books, music and movies (including downloads and streaming services), delivery charges, communication services, and education – See Figure 4	Currently manual collection using web browsers, looking to automate.
	PriceStats	Most goods, as shown in Figure 3	Working though possible sites to add coverage.
	Likely to be accessible online	Restaurant meals, tobacco, rent, transport, communication, recreation	Efficiency of collection will likely vary depending on coverage per website.
Administrative data	Government	Rent	Research underway to access coverage, methods, and timing – see Case Study 4.
	Private sector	Insurance, communication, utilities	Access likely to be more difficult than Government admin data.
Big data assisted simple survey	Government	Vehicles	Plan to simplify survey using car registration data – see Case Study 3.
Survey	Service providers	Trade services, Education, Hospital services	Surveys will likely be need for some products for the foreseeable future.

Online prices are currently included in the CPI for a growing list of products (see Figure 4). These are currently manually collected using web browsers, but we are seeking to automate this process using web robots. Many of these products are priced online to reflect consumer spending online, rather than for efficiency of price collection.

Figure 4

Current data sources

Field: physically visiting shops

Online: using web browsers

Survey: directly from businesses

Transaction: point-of-purchase scanner data

Source: Stats NZ

In addition to the online data obtained so far in the PriceStats trial, we are working through possible sites to add coverage. If we were to expand the current online data collection (PriceStats trial and operational) to try to maximise product coverage we would probably be able to cover several further significant categories. It will likely get more expensive per price observation to increase coverage. The sites chosen in the trial have on average over 10,000 price quotes per day. To obtain a similar number of price quotes for, say, restaurant meals will likely require many websites with potentially separate set-up costs for each.

Administrative data held by Government agencies and private sector companies is another potential data source. Under current cultural and legislative arrangements, obtaining access to Government administrative data is likely to be considerably easier. A significant administrative dataset for inflation measurement is the Tenancy Bond data which is currently being investigated to see if it can replace a survey based approach for rent. We describe this research further in Case Study 4. Another innovative use of Government data is our approach to simplify a survey of second-hand (used) car prices. Currently a sample survey collects prices and characteristics (for estimation using a hedonic regression model). Under the new approach, described in Case Study 3, survey respondents will only need to record prices and number plates. The number plate will be linked to administrative data to provide the characteristics needed for quality adjustment via a hedonic method.

Place dimension: Regional coverage

Turning our attention to regional coverage (a key element of the place dimension of our ideal data), some big data may have much better regional coverage than our current sample, such as transaction data for nationwide retailers or central government administrative data. On the other hand, some big data sources may lack information in this dimension. Online data is unlikely to explicitly show regional variation in prices. We will need to make a judgement about the significance of this.

Current regional price collection, where prices are collected in 12 towns and cities throughout New Zealand, is probably more grounded in history than contemporary requirements to capture price change variability. Historically, regional variation appears to have been a substantial consideration. When regular national price collection began in 1914, prices were obtained in 25 locations - a sizeable increase from the 4 main urban areas used prior to 1914 (New Zealand Government, 1915):

> The retail prices are … more liable to local, though not to temporary, fluctuations, it is advisable to extend the collection of the data as widely as possible.

Nowadays there appears to be little regional variation in price change that is significant for CPI measurement (Figure 5 shows annual price change for the 9 published New Zealand Household Expenditure Classification groups). Several recent studies have found little impact on the national CPI of excluding the price quotes obtained in the smallest regions (see Bentley, 2013). Some regional variation is apparent for lower weighted expenditure groups, but this may simply reflect random sampling variation (there is no constraining of product samples to be consistent across regions).

Housing and household utilities is a notable exception, displaying some marked differences.[6] This group accounts for about a quarter of the weight of the CPI. Fortunately, a potential big data source for rent – a major component of this group – is government administrative data which does have regional coverage. The potential to use this data for inflation measurement is explored in Case Study 4.

---

[6] Canterbury showed a sharper increase in prices, during 2012–13, following supply storages resulting from sizeable earthquakes in the region. Auckland experienced greater inflation for the group, in 2015, coinciding with steeper residential property price increases.

Figure 5

Little significant regional variation
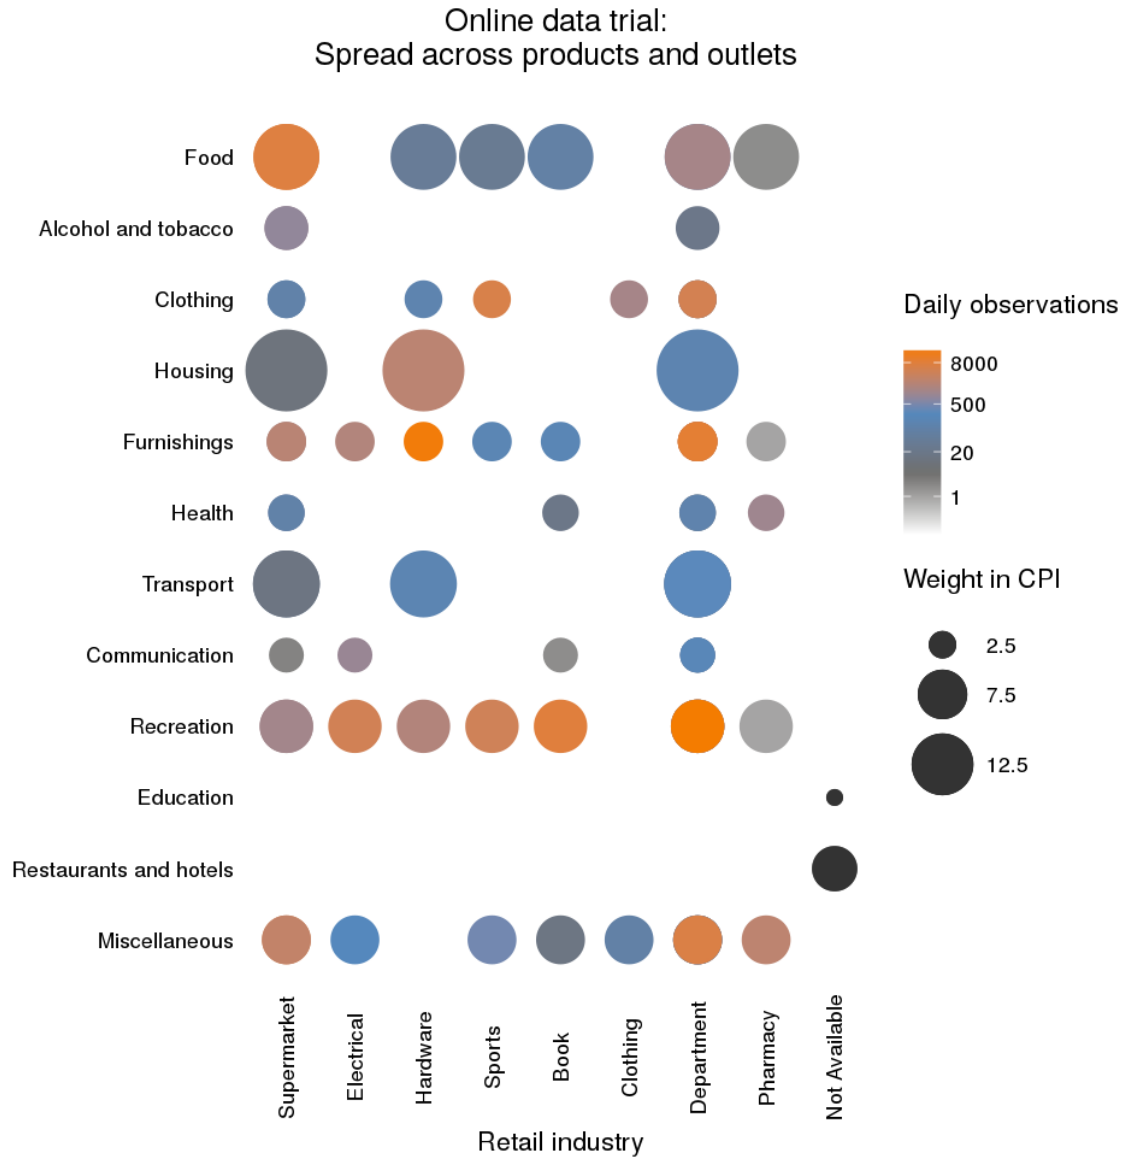

Source: Stats NZ

Place dimension: Outlet type coverage

A second consideration of the place dimension is coverage of outlet types and service providers. These capture many intangible price determining quality characteristics, such as service levels, convenience, ambience, and brand.

Figure 6 is another look at the coverage of the online trial data. The average number of daily price quotes obtained during August – December 2016, the first 4 months of the trial, is shown for retail industries by COICOP expenditure groups.[7] The chart is a reflection of a big data mind-set. We highlight the data that has been obtained, and the coverage that is achieved, then consider the significance of missing coverage. An alternative approach of seeing how well the online data covers the current data collection may show less promise as it would not reveal the supplementary data obtained.

Figure 6

---

[7] A more detailed version of this visualisation, at level 2 COICOP, is shown in the Appendix B.

Online data trial:
Spread across products and outlets

Source: Stats NZ | PriceStats

The significance of missing outlets is considered in Figure 7 which shows the most important outlet types in the Household Economic Survey, for expenditure where outlet type is recorded in the survey. We achieve coverage in industries that account for about half of all expenditure on products where the outlet type is recorded.

We didn't select any retailers to cover the predominant products in fuel, catering (cafes and restaurants) and motor trade. So these industries don't have coverage in the online data trial, but we could obtain coverage by selecting some additional retailers. Coverage will be harder to obtain for convenience stores, non-store (which includes outdoor markets) and private retailing. Coverage for specialised stores, such as liquor, meat, fruit & veg, and bakery is possible, but likely to be at a higher cost per price quote due to the set up costs to collect prices from many different websites.

Figure 7

Online data trial:
Coverage of top 20 outlet types

Proportion of expenditure (where outlet type is identified), percent
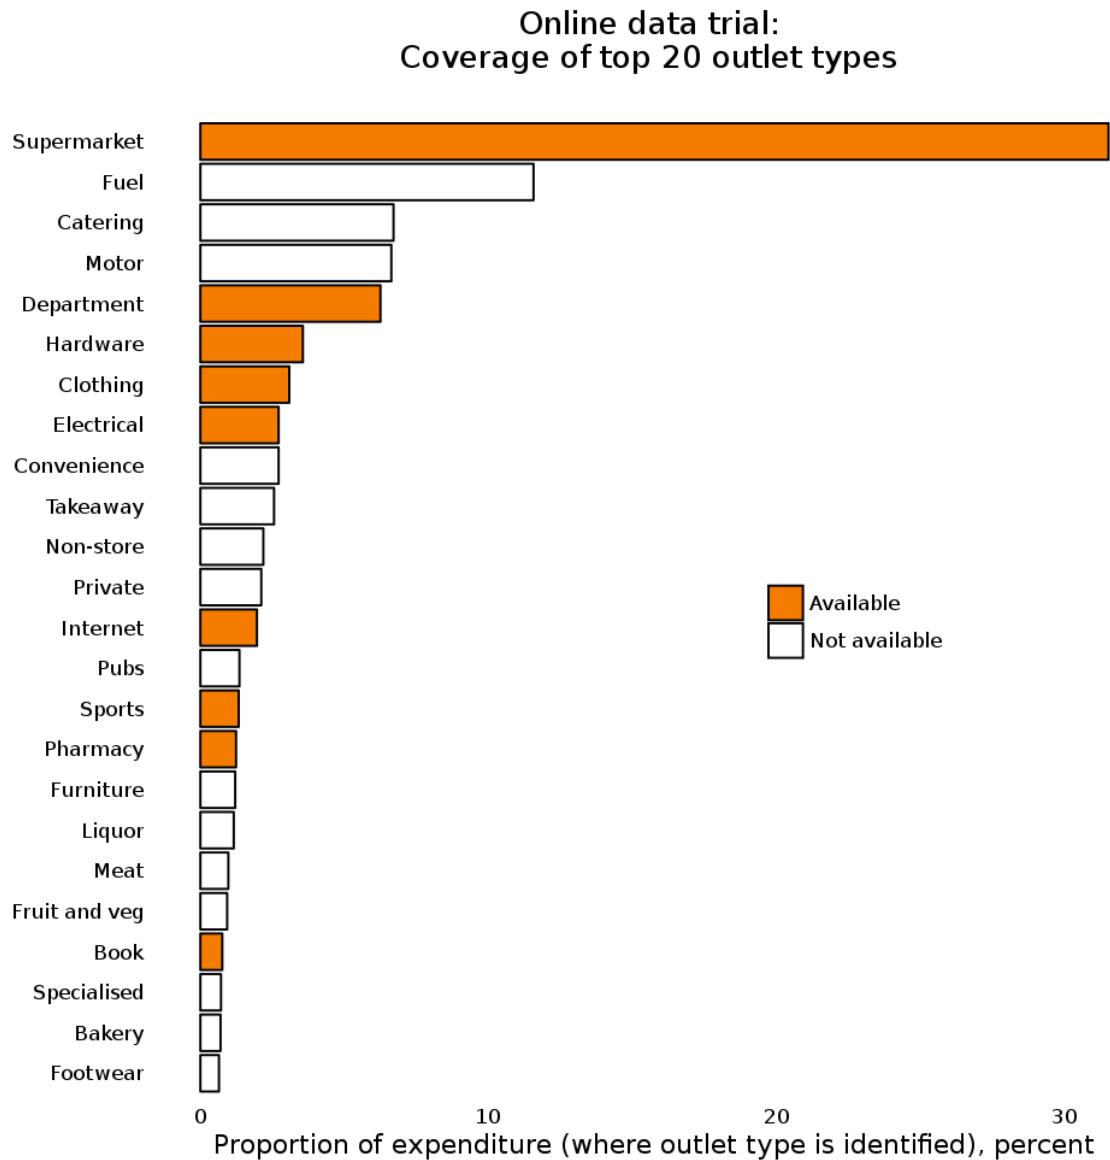
Source: Stats NZ | PriceStats

Place dimension: Online and offline prices

Online prices can be considered from two angles. Firstly, they can be considered as part of the target population of transactions. That is, to represent where consumers shop. In this sense online prices are another outlet type (with a regional dimension in the cloud). Of greater interest from a big data view point is whether or not it is appropriate to extrapolate the online prices to be representative of offline prices (such as physical stores).

Cavallo (2017) compared online and offline prices for 56 large retailers who have both online and physical stores, for 16 months starting in December 2014. The study covered retailers in 10 countries, including Australia which has many of the same retailers as New Zealand. 72 percent of prices were found to be identical across all matched observations (74 percent for the four Australian retailers, where much of the difference between channels appeared to be due to a consistent 5 percent mark-up online). Investigation into consistency in price change, based on more limited data, found "while price changes are not synchronised, they have similar frequencies and sizes". Importantly, the conclusions were positive on the potential to use online prices for official inflation measurement:

> For NSOs [National Statistical Offices] considering the use of online data for consumer price indexes, my results show that the web can be used as an alternative data-collection technology for multi-channel retailers … from a data-collection perspective, my results suggest that the online-offline price differences should not be a major source of concern.

We plan to conduct an assessment of differences, if any, between online and offline prices for some of the stores in the online data trial. This will provide a guide to the appropriateness of extrapolating online prices to be representative of offline prices.

Making an assessment of coverage will ultimately be a judgement call as we don't have perfect data about the target population of all transactions. This is not new to inflation measurement. Judgement sampling has long been used to design the sampling strategy for price collection both in New Zealand and internationally. In this sense, adopting a big-data mindset may be easier for the CPI than for official statistics that currently use random probability samples from high-quality sampling frames.

Adopting big data methods

Implementation approach: Pilot projects

We are still in an assessment phase for many potential big data sources, and we are taking a learning-by-doing approach. Our plan to adopt big data is iterative. We have already implemented transaction data in production of the New Zealand CPI for consumer electronics products, and in the New Zealand Overseas Trade Index for mobile phones and televisions. We are well positioned to start using administrative data to simplify price collection for second-hand cars. And research is underway to assess the viability of using government administrative data for rent. These examples are described in our Case Studies.

In the short-term, we are looking to assess the feasibility of replacing some of the prices currently collected in the field with webscraped data, and incorporating these using the existing price index methods of the New Zealand CPI. We also plan to automate collection of web prices that are currently collected for inclusion in the CPI. Our longer-term goal is to replace as much data as possible with digital data. To fully reap the benefits of this big data approach, model-based methods will be needed. These are described in the following section.

To work towards these goals a medium-term objective is to seek to replace our monthly food price index with big data and a model-based approach. Coverage of food prices by online data appears reasonable in the online trial data. Regional variability in food price change does not appear significant. We anticipate that the FEWS index method will work well for food prices from online data. The approach will be tested empirically by producing a big data food price index in parallel with the traditional approach. Depending on the success of this we will seek to expand this approach to other expenditure groups. A likely next step for the New Zealand CPI is that web scraped data may be able to provide more timely, 'flash estimates' of inflation that can be benchmarked to transaction (or traditional) data which may be of higher quality but slower to become available.

Model-based methods

The availability of Prices big data such as scanner data, webscraped online data or administrative data give us the opportunity to clarify what our ultimate estimation goals are.

While traditional index methods, based on a fixed basket, might be an effective (though not perfect) route to the measurement of pure price change in the context of limited data, this is not necessarily the case when we have more complete data.

Rather than using big data to 'plug into' the fixed-basket index, we try to find the optimal route to measuring pure price change, while controlling for the changing composition of the products being sold.

If we have full coverage of the product-group (eg computers, used cars, bread) that we're trying to estimate price change for, it makes sense to use all the information, rather than restrict ourselves to a fixed basket. This way we're about to properly reflect substitution across products/specs in real-time (ie when we have expenditure information as in scanner data). Also, if there is a life-cycle effect – ie a price trajectory within the life-cycle of the product – a fixed basket approach will be biased by introducing replacement products at different points in their lifecycle than the exiting products which they're replacing. Having full coverage data which captures prices as soon as products enter the market mitigates against this problem.

On the other hand, the human judgement involved in selecting high-expenditure or representative products is an advantage of the traditional approach that price measurement from online data can't replicate. Scanner data can deal with this appropriately because we have expenditures available for the weighting of the regression models. By comparing unweighted indexes from online data to weighted indexes from scanner data for the same products (Krsinich, 2015b) we can make a judgement about how well the unweighted indexes approximate the weighted indexes and therefore whether they should be used as indicative 'flash estimates' calibrated between less-timely and less-frequent higher quality indexes from weighted data, or whether they good approximations to our target price measurements.

Case studies

Case Study 1. Using scanner data for consumer electronics products in the Consumer Price Index

Status: Implemented

Data: Transaction data (point-of-purchase scanner data)

Method: Model-based (ITRYGEKS)

Stats NZ began using scanner data for consumer electronics products in the New Zealand CPI in the September 2014 quarter (Krsinich, 2014b; 2015a).

Monthly-aggregated scanner data is purchased from market research company GfK for twelve consumer electronics product categories including televisions, computers, and cellphone handsets.

The development of production processes has been an iterative one which is still evolving. This has enabled unanticipated issues to be dealt with flexibly and efficiently.

Index methodology

To produce price indexes from scanner data for consumer electronics products, we are using the Imputation Törnqvist rolling year GEKS (ITRYGEKS) index (de Haan & Krsinich, 2014).

The model-based ITRYGEKS extends the rolling year GEKS (RYGEKS) index of Ivancic, Diewert and Fox (2011) by using time-dummy hedonic (bilateral) indexes to incorporate the price movements associated with products entering and leaving the sample – ie unmatched products. GfK scanner data is valuable in that it includes very extensive information on product characteristics, so there is little risk of omitted-variables bias in the time-dummy hedonic indexes.

Practical issues associated with implementation

Quarterly indexes from monthly data

We receive monthly scanner data from GfK, but the New Zealand CPI is a quarterly index. We can either derive the quarterly index from a monthly index, or pre-aggregate the data to a quarterly level before deriving a quarterly index.

While it is useful to produce a monthly index as part of the monitoring and analysis process, it is conceptually more appropriate to derive the quarterly index from quarterly average prices and expenditure shares. This ensures the prices for products sold in each month of the quarter are appropriately weighted for price deflation, to produce quarterly volume indexes in the National Accounts.

Incomplete data at time of production

Consumer electronics data for just the first two months of the quarter are available in time to incorporate into the CPI.

There were four options for how to deal with this limitation in production. Using back-data, we assessed each of these options against the benchmark of the index we could calculate if all three months of the quarter were available.

The four options were to:

1. base the published index for the most recent quarter on the first two months of the quarter, with complete back-data feeding into the estimation (ie the third month of the quarter will be incorporated into the five-quarter estimation window for the following quarter's index calculation)
2. base the published index on three months of data, lagged by one month
3. base the published index on only the middle month of each quarter
4. base the published index on only the first two months of each quarter.

Option 1 performed best of all four options, and sits very close to the benchmark index.

Therefore, we derived the ITRYGEKS index for the latest quarter from the first two months of that latest quarter, with the third month of the quarter then being incorporated into the five-quarter estimation window used to calculate the following quarter's CPI, as shown in Figure 8.

Figure 8.

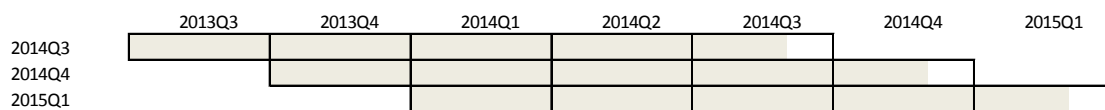**Data incorporated into each quarter's index at the time of production**

Figure 8 shows that, for example, to calculate the quarterly price movement for the third quarter of 2014 (ie the top row), we use full quarterly data from each of the quarters from quarter three of 2013 through to quarter two of 2014, and the first two months of data from the third quarter of 2014.

Developing a production process

A major aspect of using scanner data in production was developing the monitoring and analysis processes required to assure us of the data quality, and to give context to the indexes.

These processes include:

- Checks that sets of characteristics have remained the same since last quarter
- Time series (across the full 5 quarter estimation window) of monthly quantities sold for each product category
- Time series of monthly average prices (unadjusted for quality change) for each product category
- Identification of outliers, in terms of both movements and levels for the latest quarter, and incorporating information from the full longitudinal record.
- Time series of distributions across key characteristics, for each product category, in terms of both quantities sold and expenditure shares
- Monthly ITRYGEKS quality-adjusted price indexes

- Time series of turnover statistics – ie counts of new products, disappearing products, and matched products between each pair of adjacent months

If one of the monitoring or analysis outputs suggests further investigation, other processes can then be invoked:

- The quarterly (or monthly) price indexes can be run for subsets of the data. For example, we could run the quarterly ITRYGEKS for different brands of television, to determine whether an unusual price movement is contributed to by all brands, or whether one particular brand is driving the result.
- We can run the price indexes excluding ouliers – either full longitudinal records of outlying products, or particular outlying values of products. So, if outlier detection process identifies a potential error, we can test the impact on the index of outlier removal.

Case Study 2. Import data for mobile phones and televisions in the NZ import price index

Status: Implemented

Data: Government administrative

Method: Model-based (FE)

In the December 2013 quarter, Stats NZ started using administrative data for mobile phones and televisions in the New Zealand import price index.

For the major brands of these two products comprehensive data on import prices and quantities is available at a detailed product level.

Index methodology

We use a fixed-effects (FE)[8] index method to estimate the price indexes. Data is available at the detailed product level, so we control for compositional change by including the product identifier in the hedonic regression models. Krsinich (2016) shows that this is equivalent to a fully-interacted time-dummy hedonic index, when all characteristics are expressed as categorical variables.

Because the previous quarter of the New Zealand import price index is revisable, it is not necessary to apply the window-splicing aspect of the FEWS aspect to avoid systematic bias due to missing the implicit price movements of new products entering the market. However, window-splicing may be incorporated in the future to improve the estimation by incorporating updated estimates of fixed-effects parameters as more data is incorporated for new products.

Monitoring

The data is supplied to us in different formats for the different brands of phones and televisions. These are made consistent across brands and across time.

We produce the indexes at the brand level, to give context to the aggregate price movements.

An important validation step is to overlay the index ending in the current quarter, over the (rebased) index from last quarter's production, to check that there is little 'window drift' with the updating of the estimation window. Any problems with the most recent quarter's data not already identified are likely to emerge at this stage.

---

[8] Also known as a time-product dummy (TPD) index (see Ivancic, Diewert & Fox, 2011)

Case Study 3. Used cars

Status: Model Implemented, researching the incorporation of administrative data

Data: Big-data assisted simple survey

Method: Model-based (time-dummy hedonic index)

Each quarter, data on the sales of approximately 3,500 cars are collected from a sample of used car dealers. Price, year of manufacture, make and model, engine size (cc rating), and odometer reading are collected for each car sold.

The sample was designed initially to support the calculation of average prices within estimation cells based on combinations of make and model, cc rating ranges, and age of car. These were then weighted together, and the price index derived from the aggregate average prices in each quarter. To ensure robust estimation of averages, cells with too few observations were excluded, which resulted in the use of only around 25 percent of the data collected.

Initial hedonic method for used cars

In 2001, Stats NZ began using a hedonic approach to more efficiently make use of the collected data. This initial hedonic estimation was a regression of price (of car) against time dummies and the observed characteristics of the cars. In these early days when the hedonic approach was relatively new to us we were able to justify and explain this linear formulation of the hedonic model as a more efficient version of the existing estimation method, using the full sample of data and therefore being more representative of the target population:

$$P_c = \sum_k \beta_k C_k + \sum_t \delta_t D_{ct} + \varepsilon_{ct} \qquad (1)$$

Where:

P_c is the price of car c (note that there is no t term as we can assume each individual car is sold only once with the same characteristics)

$D_{ct} = 1$ in the quarter t that car c is sold, and 0 otherwise

The following characteristics K were included in the initial hedonic model for each car c sold:

- city/town of purchase (15 categories)
- make and model (47 categories)
- age of car (in years)
- size of engine (cc rating, eg 2300)
- odometer reading (in kilometres).

The updated hedonic model for used cars

In 2011 we improved the hedonic model by fitting log of price, rather than price, adding extra characteristics, and adding squared terms for age of car and size of engine (Krsinich, 2011).

The updated model now used in production is:

$$\ln P_c = \sum_k \beta_k C_{kc} + \sum_t \delta_t D_{ct} + \varepsilon_{ct} \qquad\qquad (2)$$

With P_c and D_{ct} as for (1). Characteristics K are now:

- city/town of purchase (15 categories)
- make and model (96 categories)
- age (in years) and age squared
- size of engine and size of engine squared
- odometer reading
- dealer (approximately 300).

To operationalise the hedonic method in production, a rolling window of the latest eight quarters is used to estimate the hedonic model each quarter. The derived index for the most recent quarter is linked to the previous quarter's index number. In other words, we are using a 'movement splice' (de Haan, 2015) to update the unrevisable index.

The used cars index was Stats NZ's first production use of hedonic methods, so we built a process of monitoring and analysis around that which has informed the later use of big data (in particular the consumer electronics scanner data). In particular, the comparison of the most recently estimated index across the full 2-year estimation window to the rebased 2-year index ending in the previous quarter helps to assure us that the use of a movement splice is a valid approach (ie because the indexes don't drift apart over time).

Incorporating administrative data and reducing sample size

We now plan to incorporate New Zealand Transport Authority (NZTA) data into the survey, by asking the used car dealers to give us just the registration numbers of cars sold – for which we'll then link on characteristics. This reduces the respondent burden while making use of existing administrative data.

At the same time we plan to reduce the used car dealer sample size. The sample was originally designed to achieve desired accuracies for the estimation cell approach. The hedonic index is a much more efficient use of the data, so there is scope for sample size reduction.

Case Study 4. Rent

Status: Preliminary research

Data: Government administrative

Method: Model-based (likely to be FEWS method)

Rent (actual rentals for housing) is one of the most important components of the CPI. It is about 10 percent of the CPI by expenditure weight. For households who pay rent, the proportion of their expenditure on rent is typically 30–40 percent. We have begun a research project to assess the potential to replace our current sample-survey approach with big data – Tenancy Bond administrative data, held by the Ministry of Business, Innovation and Employment (MBIE). This case study presents our initial investigations.

Current data collection is a postal survey of landlords.[9] The survey population is all identified rented dwellings within the sampled geospatial areas. A matched-sample approach is used to control for the changing quality of the stock of rental dwellings (Krsinich, 2011).

The Tenancy Bond data is a digital footprint created as a by-product of an administrative process. Landlords can ask tenants to pay a monetary bond as security when they move into a property. Landlords who charge a bond must lodge it with MBIE's Tenancy Services within 23 working days (Tenancy Services, 2017). The Bond lodgement form (which can be completed online or by post) includes a requirement to state the weekly rent payment. Other data captured includes the dwelling address, dwelling type (such as room, flat, house), and the number of bedrooms. A unique property ID is created as part of the administrative process.

Beyond its significance in the CPI, rent statistics are of considerable public interest. The CPI rent series, however, are not always the primary source of information. Limiting factors include the relatively infrequent, quarterly, publication and lack of regional detail. Statistical power is small at a subnational level as a result of a nationally-designed sample survey of about 2,000 dwellings. Changes in median rent prices directly from MBIE's tenancy bond data are often quoted. As are rent price changes from the online marketplace and advertiser, TradeMe. Both these data sources provide monthly updates and a greater regional breakdown. Open Data provided by MBIE is substantially more detailed. What these data sources don't account for is the changing quality of rental properties.
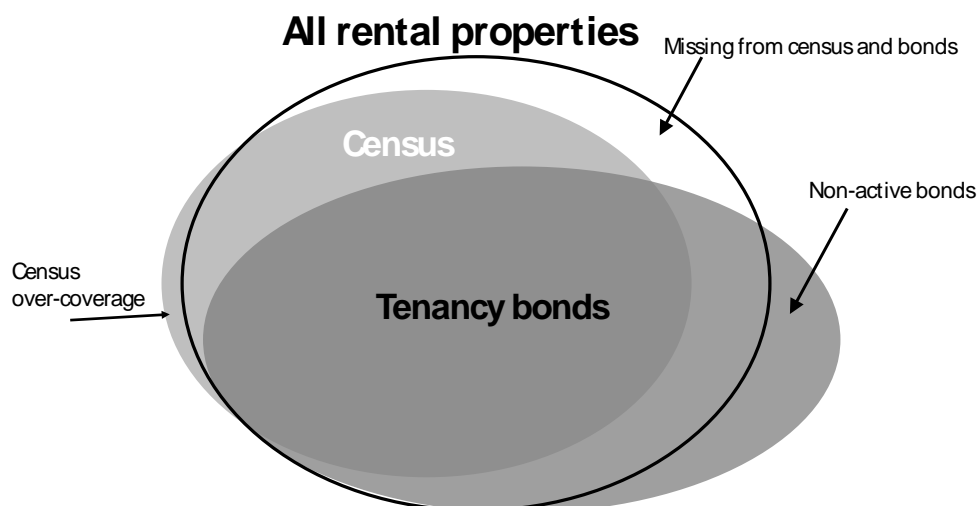
Coverage

The target population is expenditure on all private rental dwellings by New Zealand resident households. There is no comprehensive dataset listing all in-scope rental properties, yet we can compare the leading data sources against each another to assess their strengths and weaknesses. Figure 9 is a stylised visualisation of the coverage of the current gold standard, the New Zealand Census of Population and Dwellings, compared with the Tenancy Bond data. From dwellings that report an amount for the rent paid question, we can estimate the total stock of rental dwellings and rent price statistics for these. This census-derived population is used as the sampling frame for the current rent postal survey.

---

[9] The survey covers private residential dwellings. Inflation for government-owned (Housing New Zealand) properties is calculated from administrative data.

Figure 9

Coverage of tenancy bond administrative rent data

The detailed cross-sectional regional and demographic information in the census is at the expense of temporal coverage. The snap-shot nature of the census, focusing on dwellings as at census night, means dwellings for households who are temporarily away from home on census night would not be included in the census rent-paid-dwellings data. The impact of this, along with the impacts of other coverage issues (described in table 2), can be explored by comparison with population estimates of the stock of rental dwellings. This data is based on extrapolation of owner-occupier rates, from a home ownership question in the census, applied to estimates of the total New Zealand private dwelling stock. This latter approach assumes that unoccupied dwellings on census night are rented and owned in the same proportion as those occupied. This may not be true.

Table 2
Coverage limitations: Data source comparisons

Data source	Possible under-coverage	Possible over-coverage
Census	• Dwelling non-response • Item non-response to rent-paid question • Unoccupied on census night	• Intra-household rent payments, for example child to parent • Non-resident household expenditure
Tenancy Bond	• Informal arrangement, where landlord does not require bond • Illegal tenancies, where bond has not been lodged	• Unclaimed bonds • Non-private households
Population estimates	• Extrapolated owner-occupier rates	• Extrapolated owner-occupier rates

Known coverage limitations of the tenancy bond data are the exclusion of rentals with no bond lodgement form, and unclaimed bonds after a tenancy has ended. Lack of lodgement includes

rentals where the landlord has not asked the tenants for a bond, and any bonds that are not compliant with the law. Active bonds are identified in the administrative data based on bond lodgements and returns of bonds. Unclaimed bonds represent over-coverage.

Figure 10

Growing coverage of tenancy bond data
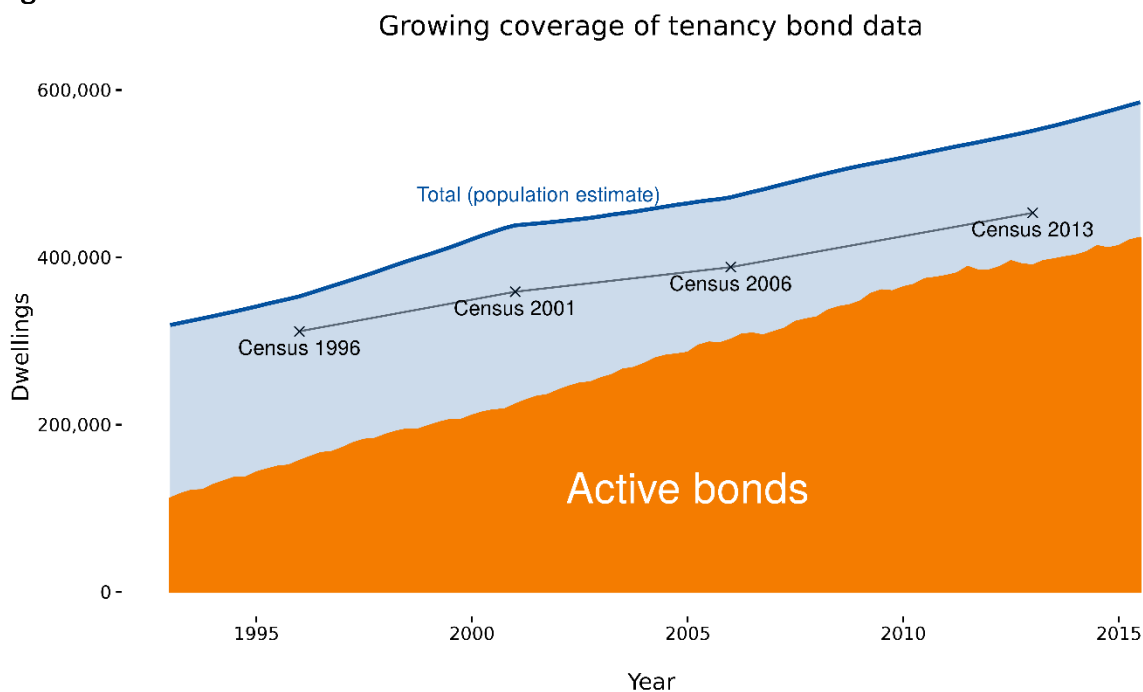
Source: Stats NZ | Ministry of Business, Innovation and Employment

Coverage of the tenancy bond data has increased over time as new tenancies were progressively required to become compliant. We show the increasing coverage over time in figure 10. At the time of the most recent census, in March 2013, there were about 390,000 rental dwellings with active bonds compared with about 450,000 rental dwellings in the census and 550,000 rental dwellings using population statistics estimates. This suggests the tenancy bond data covers 86 percent of the rental dwellings in the census and 70 percent of all New Zealand rental dwellings.

- Distributional analysis of coverage

For inflation measurement, we only need representative coverage, rather than exhaustive coverage for national aggregates. This makes distributional comparisons invaluable to help understand if the missing rentals are likely to be missing at random (or not). Figure 11 shows the regional distribution of bond coverage of the census.[10] We find better coverage in the bigger regions (perhaps, suggesting more informal rental agreements in smaller settlements).[11]

Suei (2016) compared the rent amount recorded in the Tenancy Bond data with the rent amount reported in the Census of Population and Dwellings, at both aggregate and dwelling

---

[10] Population statistics estimates of rental dwellings is not available regionally.

[11] The comparison also shows more rental properties in the Bond data for Otago than are present in the census. This region has a large university student population. The apparent over-coverage may represent unclaimed bonds or unoccupied dwellings on census night, or census non-response.

levels.[12] The study found near identical distributions for weekly rent amount, number of bedrooms, and sector of landlord.

Figure 11

Better coverage of main centres
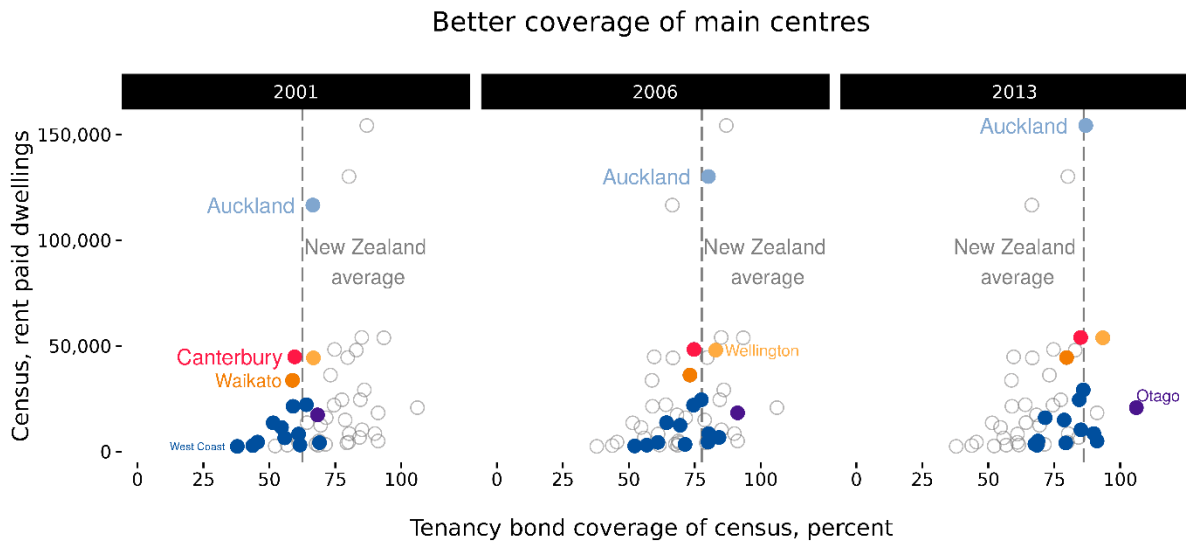
Source: Stats NZ | Ministry of Business, Innovation and Employment

Another way to assess the coverage of the bond data is to look at the new price information that is available each quarter. In a sense, this is the sample size (sampling ratio) of the quarterly price information. Figure 12 shows the ratio of lodged bonds to active bonds, in a given quarter. The ratio has been decreasing. This reflects the increased coverage of active bonds as well as, potentially, changes in the average duration of tenancies. The proportion of quarterly lodged bonds appears to be stabilising at around 10 percent. We are planning to compare this rate with the survey birth and death rates, and proportion of surveyed dwellings showing price increases and decreases. Using the microdata we can also calculate average tenancy lengths. In a static population, a 10 percent coverage rate would indicate average tenancies of about 2.5 years.

---

[12] The microdata was linked in Stats NZ's Integrated Data Infrastructure (IDI) using address-matching. About 45 percent of census dwellings were successfully matched to Tenancy Bond data.

27

Figure 12

Stabilising coverage of lodged bonds
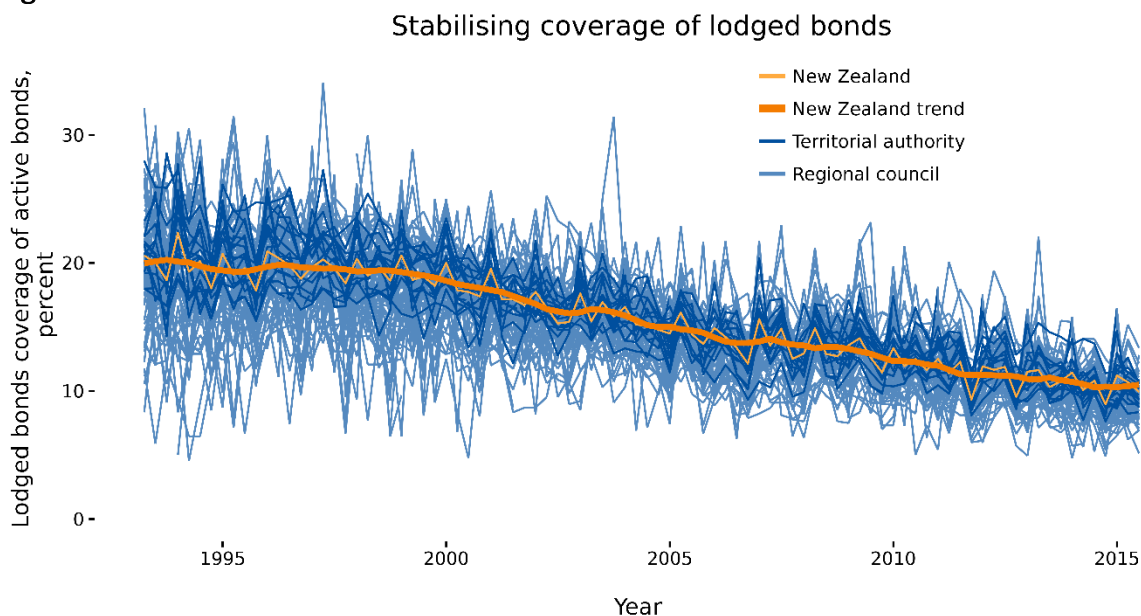
Source: Ministry of Business, Innovation and Employment

Overall, the coverage of the tenancy bond data appears reasonable. There are limitations with all data sources, but coverage issues identified with the administrative data do not appear prohibitive.

Quality adjustment

Current publication of tenancy bond data by MBIE does not include a measure that controls for the changing mix of dwellings overtime to measure pure price change. Such a method is required to use the data to measure rent price inflation, for inclusion in the CPI. The bond data has limited explicit quality characteristics, therefore our preferred approach for investigation is FEWS. Krsinich (2011, 2016) provides the justification for using fixed effects for rent inflation measurement, where we lack full information on price-determining characteristics of rental dwellings. We plan to empirically test the approach. This will include sensitivity analysis to help determine an appropriate window length.

Timing of recording price change

The final issue that needs to be addressed is to consider the timing of recording price change. Rent price changes are only observed in the bond data when tenancies begin, or a new bond lodgement form is submitted (for example, if the landlord wants to increase the bond amount as well as the weekly rent). In contrast, price changes are observed in the current, survey-sample approach whenever the rent amount changes, regardless of whether this is a new or existing tenancy. To adopt the digital data source we need to make a judgement about the significance of the difference in timing of recorded price change. We consider this from two perspectives:

- How well does the admin data proxy the current approach?

- What price change would we ideally want to reflect in the CPI?

Suei (2016) compared the weekly rent amount recorded in the Tenancy Bond data with the rent amount reported in the census. For dwellings that could be matched between data sources, 89

percent were in the same rent amount band.[13] For those not in the same band the vast majority had higher rent recorded in the census. For point-in-time estimates (which was the focus of Suei's study), this reflected the out-of-date nature of some tenancy bond data. The administrative data is a reasonable overall proxy to the current approach, but there is a noticeable timing difference.

To look at the timing from a different perspective is another reflection of our big data mind-set: What price change would we ideally want to reflect in the CPI? How closely does the big data meet this need? This is a paradigm shift from simply considering how well the big data matches Stats NZ's current approach. Figure 13 shows that using the change in the average price of newly lodged bonds looks to be a leading indicator of rent price change. This is somewhat counter-intuitive to what we might expect from the point-in-time conclusion that some of the tenancy bond price information is out-of-date. The critical difference is that using the price information from newly lodged bonds reflects the current market price for rental properties. Taking this one stage further, the within tenancy rent price changes (captured in the current survey, but omitted from the administrative data) perhaps don't reflect true market rent. They may reflect a discount for loyalty and cost savings to the landlord compared with finding replacement tenants.

Figure 13

Tenancy bond data appears a leading indicator, but does not account for quality change
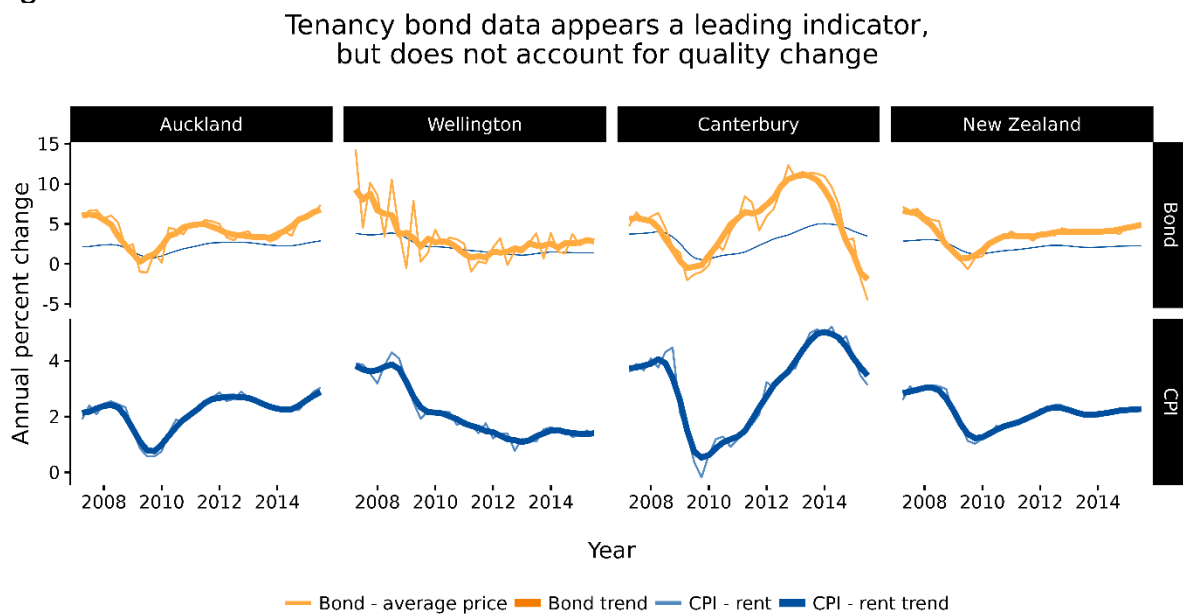
Source: Stats NZ | Ministry of Business, Innovation and Employment

Assessment

Preliminary assessment of the coverage of the big data is looking promising. We still need to determine a suitable quality adjustment method to control for the changing mix of dwelling characteristics to measure 'pure' price change. This will likely be a fixed effects model (with dummy variables for individual dwellings). Further, we need to determine the suitability of the observed price changes in the bond data – that is, price changes only at the change of tenancies – and make a judgement about the significance of this.

---

[13] Rent amount bands used were: Under $50; $50-$79; $80-$99; $100-$124; $125-$149; $150-$174; $175-$199; $200-$249; $250-$299; $300-$349; $350-$399; $400-$499; $500+

Concluding remarks

A little think about big data and we begin to wonder: are we making the most of today's digital data for measuring inflation? Why are we physically visiting stores to collect prices when these are available online? Can we really justify the costs and respondent burden of all our postal surveys? Can't we do something more efficient than manually collecting prices using web browsers? Does our current approach fit with our modern public service? Are we best meeting customer needs?

These questions lead to a need for a bigger think about big data for inflation measurement. This paper is a contribution to that thinking. To maximise the use of big data we need a mind-set shift towards accessing the coverage of digital data and using model-based approaches to inflation measurement. Competing sources of price information necessitate a need to critically assess these and seek to adopt, adapt or explain. Big data may not always provide the information needed, but in these cases we need to be able to clearly articulate why not.

References

Adobe (2017). Adobe Digital Price Index. Retrieved from
https://blogs.adobe.com/digitalmarketing

ANZ (2017). Monthly Inflation Gauge. Retrieved from www.anz.co.nz

Australian Bureau of Statistics (2016). Making greater use of transactions data to compile the consumer price index, Australia. Retrieved from www.abs.gov.au

Bentley, A (2011). Accounting for quality change in the CPI. In Asymmetric Information, Issue No. 42/December 2011. New Zealand Association of Economists, Wellington. Retrieved from www.nzae.org.nz

Bentley, A (2013). Sampling Framework. In Consumers price index advisory committee 2013 discussion paper. Retrieved from www.stats.govt.nz

Bentley, A (2014). Special consumer price indices for particular groups of households. In Forbes, S, & Victorio, A (Eds) The New Zealand CPI at 100: History and interpretation. Wellington: Victoria University Press.

CardSmart (2017). PriceWatch. Retrieved from www.pricewatch.co.nz

Cavallo, A & Rigobon, R (2016). The Billion Prices Project: Using Online Prices for Measurement and Research. *Journal of Economic Perspectives, 30*(2):151-178. Retrieved from www.mit.edu

Cavallo, A (2017). Are online and offline prices similar? Evidence from large multi-channel retailers. *American Economic Review 2017, 107*(1):283-303. Retrieved from www.aeaweb.org

de Haan, J (2015). Rolling Year Time Dummy Indexes and the Choice of Splicing Method. Room document tabled at the Ottawa Group, Tokyo, Japan. Retrieved from www.stat.go.jp

de Haan, J, & Krsinich, F (2014). Scanner data and the treatment of quality change in nonrevisable price indexes. Journal of Business and Economic Statistics, 32(3). Available from www.tandfonline.com

de Haan, J & Krsinich, F (in press). Time dummy hedonic and quality-adjusted unit value indexes: Do they really differ? Review of Income and Wealth.

Krsinich, F (2011). Measuring the Price Movements of Used Cars and Residential Rents in the New Zealand Consumers Price Index. Paper presented at the Ottawa Group, Wellington, New Zealand. Retrieved from www.stats.govt.nz

Krsinich, F (2014a). Quality adjustment in the New Zealand Consumers Price Index - Current practice and future directions. In Forbes, S, & Victorio, A (Eds) The New Zealand CPI at 100: History and interpretation. Wellington: Victoria University Press. Retrieved from www.researchgate.net

Krsinich, F (2014b). Measuring price change for consumer electronics using scanner data, Statistics NZ, Wellington. Retrieved from www.stats.govt.nz

Krsinich F (2015a). Implementation of consumer electronics scanner data in the New Zealand CPI. Paper presented at the Ottawa Group, Tokyo, Japan. Retrieved from www.researchgate.net

Krsinich F (2015b). Price indexes from online data using the fixed-effects window-splice (FEWS) index. Paper presented at the Ottawa Group, Tokyo, Japan. Retrieved from www.researchgate.net

Holt, J (1986). *Compulsory arbitration in New Zealand: The first forty years*. Auckland: Auckland University Press.

ILO, IMF, OECD, UNECE, Eurostat & the World Bank (2004). Consumer price index manual: Theory and practice. Retrieved from www.ilo.org

Ivancic, L., W.E. Diewert and K. J. Fox (2011). Scanner data, time aggregation and the construction of price indexes. *Journal of Econometrics 161*(2011):24-35.

Ministry of Business, Innovation and Employment (2017). Rental bond data. Retrieved from www.mbie.govt.nz

Metcalf, Flower, Lewis, Mayhew and Rowland (2016). Research Indices, using web scraped price data: clustering large datasets into price indices (CLIP). Office for National Statistics, UK. Retrieved from www.ons.gov.uk.

New Zealand Automobile Association (2017). Petrolwatch. Retrieved from www.aa.co.nz**Fehler! Hyperlink-Referenz ungültig.**

New Zealand Government (1915). Report on the cost of living in New Zealand, 1891–1914. Nov 1915. Wellington: Government Printer. Retrieved from www.stats.govt.nz

New Zealand Government (1920). An inquiry into Prices in New Zealand, 1891–1919. Wellington: Government Printer. Retrieved from www.stats.govt.nz

State Services Commission (2016). Better public services. Retrieved from www.ssc.govt.nz.

Statistics NZ (2006). Consumers price index review 2006. Retrieved from www.stats.govt.nz

Statistics NZ (2013). Report of the Consumers Price Index Advisory Committee 2013. Retrieved from www.stats.govt.nz

Statistics NZ (2014). Consumers price index review: 2014. Retrieved from www.stats.govt.nz

Statistics NZ (2016a). Household living-costs price indexes: Background. Retrieved from www.stats.govt.nz

Statistics NZ (2016b). Statistics New Zealand's Strategic intentions for the period 2016/17–19/20 Annual report for the year ended 30 June 2016. Retrieved from www.stats.govt.nz

Suei, S (2016). Comparing housing information from census and tenancy bond data. Unpublished report. Statistics NZ, Wellington.

Tenancy Services (2017). Bond. Retrieved from www.tenancy.govt.nz

TradeMe (2017). Rent Price Index. Retrieved from www.trademe.co.nz.

Appendix A: The century-old inflation measurement paradigm

The underlying philosophy of price measurement, and many of the practicalities and conventions, adopted at the start of the New Zealand CPI, over 100-years ago, still form the bedrock of price measurement today. This section elaborates. A focus of this paper is to challenge our mind-set to price measurement and consider whether New Zealand's current approach is optimal for today's digital-age, or overly anchored in history.

In New Zealand, regular national price collection began in 1914 (New Zealand Government, 1915). This was based on a basket of 67 sampled products in 25 towns. Publication was quarterly, and data collection was monthly for all products, except rent which was collected half-yearly. A sample of products, in a sample of stores, in a sample of urban areas, at sampled points in time.

Regional coverage was a prominent consideration – expanding price collection to 25 towns was a substantial expansion from the earlier data being limited to the four main urban areas (New Zealand Government, 1920):

> To obtain a general estimate of the course of prices for the whole Dominion [of New Zealand] it is clearly insufficient to collect data merely from the four chief centres. This was early recognized, and from the beginning of the year 1914 particulars of retail prices have been obtained for each month in twenty-five different towns of New Zealand. The twenty-five towns were selected as representative of New Zealand as a whole; they cover both Islands, from Whangarei to Invercargill, and represent coastal and inland districts and large and small centres.

Regional variability a century ago is likely to have been caused by economic factors that are unlikely to persist to the same extent today. Discussion of the regional coal market is perhaps a good example where reduced distribution costs are radically different today (of course, different forms of consumer energy would also predominate these days) (New Zealand Government, 1915):

> It will be very noticeable that Auckland and Dunedin prices of coal are very considerably lower … The explanation seems to be the use of local coals … Christchurch and Wellington, on the other hand, use the more expensive coals … These sea-borne coals are necessarily more expensive, especially when, as on the West Coast, bar harbours hinder shipping. In Christchurch, too, a heavy tunnel rate must be paid on all sea-borne goods.

In 1914, price collection was a mix of visiting stores (by the Inspector of Factories in each of the 25 towns), and direct returns from a sample of grocers, butchers, and house agents.

Product coverage was a small sample, despite an observation that "the selection of the commodities to be included is perhaps the most important step". The need to take a sample was justified on practical grounds. Noting that the "general tendency" of prices could be measured by selecting "typical commodities" that represent the various kinds of commodities in general use, the individual prices tracked were based on the "most frequent prices of the predominant brands".

Monthly prices were collected on the 15th of each month. Collecting rent prices every 6 months was justified by considering changes in rent to be relatively infrequent. Frequency of publication did not appear to be a primary concern. The main use of the early New Zealand CPI was for cost-of-living adjustments to wages, principally by the Arbitration Court (Holt, 1986). This use did not require high-frequency data as wage rates are typically adjusted no more frequently than annually.

Price index methods

Report on the cost of living in New Zealand, 1891–1914 (New Zealand Government, 1915) provides a summary of the rationale for the price index methods that have been used to calculate aggregate inflation (or rather, as was the focus at the time, changes in the 'cost-of-living'). Described as the "aggregate-expenditure method", the approach adopted is the fixed-basket, Laspeyres-type, index formula that has been used internationally as the principal approach to inflation measurement for the past century.

> The measurement of changes in prices is a problem of the utmost complexity … But, on the other hand, there is always noticeable a general tendency, around which the particular prices may oscillate, but which is capable of being measured … To measure changes in this 'general level of prices' there has evolved the method of index numbers.

Recommending this approach appeared to be largely pragmatic. Advantages such as simplicity and consistency with Australia were noted, alongside the problems of adopting a fixed-basket approach, such as assuming constant purchases over time.
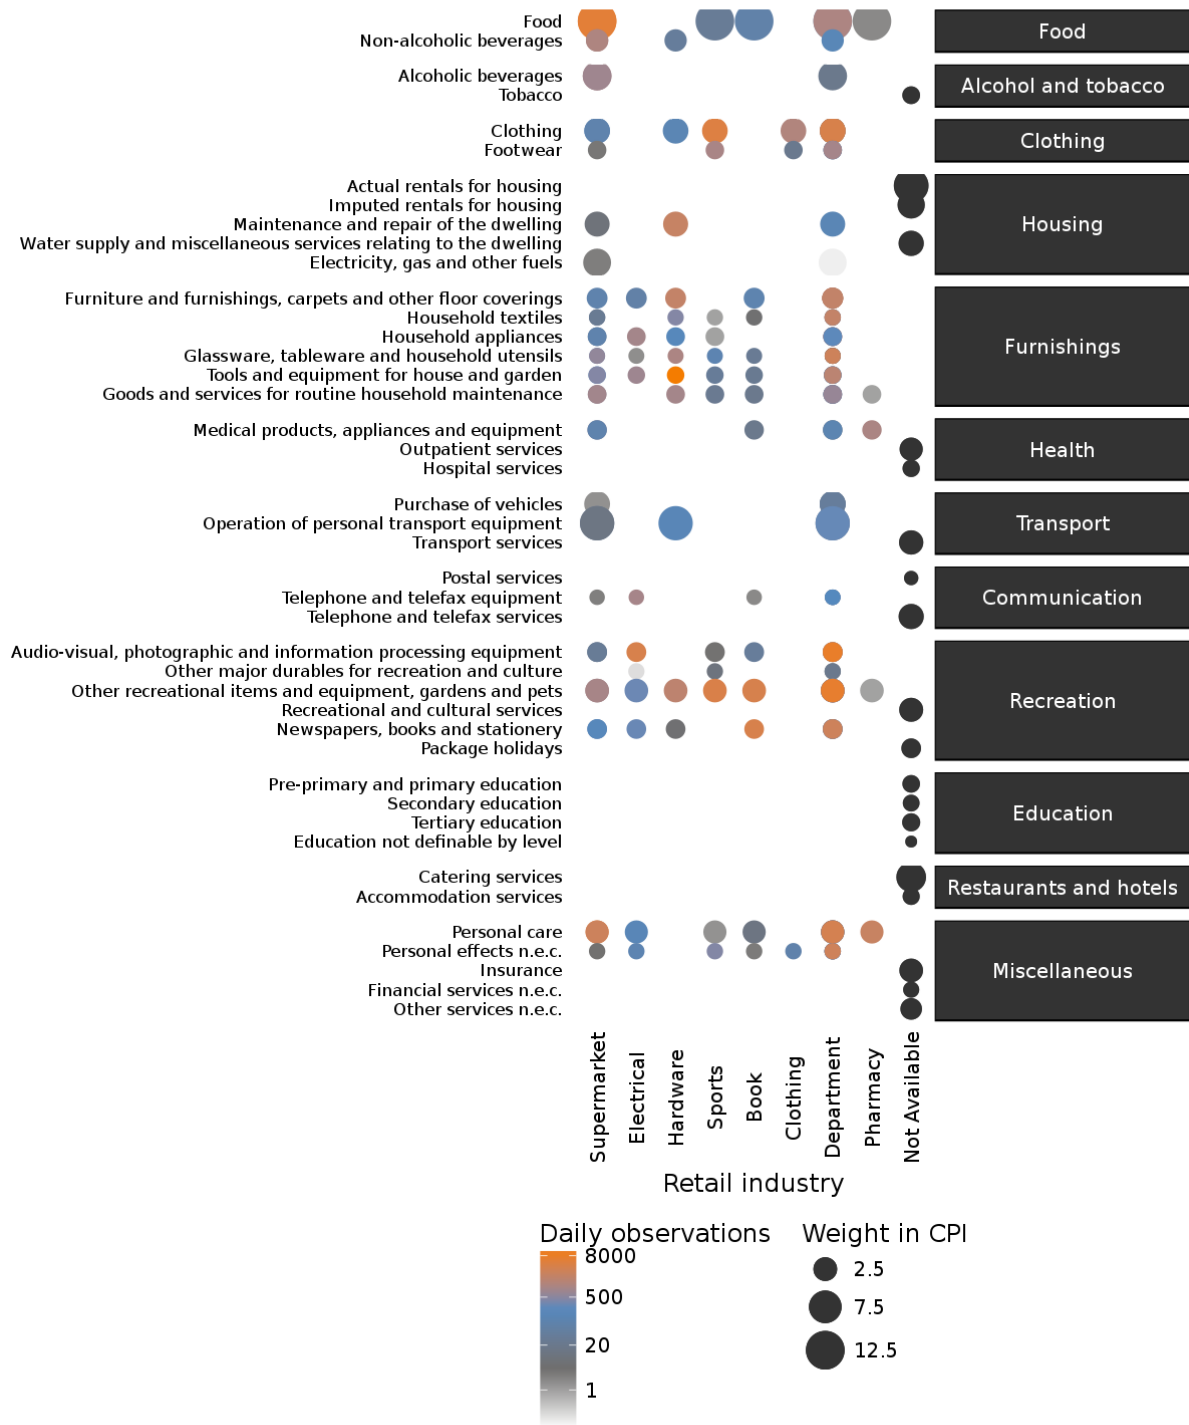
> It is admitted that the method assumes a fixed consumption of the articles treated, and over a course of years this may become a serious defect; but by this assumption it is possible to measure changes in prices very accurately … in retail prices especially, the relative consumptions or usages do not vary greatly from year to year, and new commodities are introduced but slowly.

Appendix B: Detailed coverage of the online data trial

This visualisation shows the average number of daily price quotes collected from 13 New Zealand retailers during August-December 2016, as part of our collaboration with PriceStats.

s

Online data trial:
Spread across products (2nd level COICOP) and outlets

Source: Stats NZ | PriceStats