Bigger, better, faster: further progress in using non-traditional data to measure price inflation

Paper presented at the 61st Annual Conference of the New Zealand Association of Economists,
at Wellington, New Zealand,
23-24 June 2021

Matthew Stansfield, Frances Krsinich, Prices, Stats NZ
P O Box 2922
Wellington, New Zealand
info@stats.govt.nz
www.stats.govt.nz

Citation
Stansfield, M & F Krsinich (2021, June). *Bigger, better, faster: further progress in using non-traditional data to measure price inflation* Paper presented at the 61st Annual Conference of the New Zealand Association of Economists, Wellington, New Zealand.

Abstract

In recent years Stats NZ has been collecting and using price information from a variety of different non-traditional sources such as web-scraping, point-of-sale scanning, and administrative sources such as tenancy bonds and customs.

Timely, granular, and high coverage, these present both opportunities and challenges for standard price index methods and production, in particular due to their volume and churn.

This paper details Stats NZ's investigation and plans for developing a framework to integrate and calculate quality adjusted price indexes from these non-traditional sources. We present the new approaches to estimating price movement from this kind of data and show empirical results comparing how these perform for different product classes.

Introduction

Price indexes are traditionally estimated using periodic measurements of the prices of a ‘fixed basket’ of goods, carefully chosen to represent the average consumer’s purchases. Expenditure weights are used to determine the relative weights of each product in the basket, and these are updated three-yearly based on household expenditure surveys.

With increasingly rapid change in technical innovation and product development, the fixed basket approach is gradually becoming less able to reflect pure price change in the context of both rapid quality change, and the introduction of entirely new products in the marketplace.

Alongside these escalating pressures on the traditional price index methods, there is increasing availability of rich new price data as a result of computerisation of sales, administrative data and online shopping sites. Increased processing power and memory capacity enables more computationally intensive methodologies to be implemented in research and production.

Alternative sources of price data include:

- point-of-sale data – also known as scanner or transactions data
- administrative data such as tenancy bond data, NZTA data, customs data
- web-scraped data – also known as online data

The advantage of these alternative data sources are that:

- they exist already, allowing the reduction of respondent burden and savings in field collection resource
- they often provide close-to full coverage of the relevant product universe, compared to the relatively small product samples used under the traditional framework
- some sources, such as scanner data, capture quantities as well as prices at the most detailed product specification level, allowing us to reflect product substitution in close-to real time, as it happens in response to price changes and changing consumer behaviour.

The disadvantages of these alternative data sources are that:

- although the application of traditional (bilateral) index methods is inadequate there is, as yet, no international consensus on what the best multilateral methods for this kind of data are, though opinions are starting to converge

- as with any big data, there is a risk that users won't identify problems in the data, so comprehensive diagnostic processes are required, both for the ingesting of data and for the analysis of resulting price movement estimates, to identify products/records with significant effects on movements
- the coverage of the data can diverge from the conceptually ideal. For example: supermarket scanner data includes some expenditure by businesses (conceptually out of scope of the CPI), and consumer electronics scanner data can't be disaggregated to regions due to retailer confidentiality constraints.

There has been a huge upswing in international research and development of methods and processes into utilise this kind of data for price measurement over the last couple of decades. Stats NZ has contributed to this research as an early adopter of alternative data sources such as:

- scanner data for consumer electronics products (Stats NZ, 2014)
- tenancy bond data for the rent price index (Bentley, 2018; Stats NZ, 2019a), and
- customs data for export and import prices in the Overseas Trade Index (Stansfield, 2019; Stats NZ, 2019b)

Implementation of these new data sources has used so-called 'multilateral methods'[1] such as:

- the Time Dummy Hedonic (TDH) index for used cars (Krsinich, 2011)
- the Imputation Tornqvist Rolling Year GEKS (ITRYGEKS) index for consumer electronics products (de Haan & Krsinich, 2014), and
- the fixed-effects window splice (FEWS)[2] index used, for the rent price index and the overseas trade index mentioned above (Krsinich, 2016)

Krsinich (2014), Bentley & Krsinich (2017) and UNESCAP (2019) all give overviews on the last twenty years of developments at Stats NZ in the adoption of these new data sources and the development of the corresponding new index methods.

Why new methods are needed for alternative price data

As alternative sources of price data become available, and computing power becomes sufficient to deal with them, applying traditional index calculation methods becomes both unrealistic and insufficient. The new data sources are generally characterised by large volumes and high product churn – i.e. introduction and disappearance of products. It is therefore necessary to use methods that can 'automatically' handle quality change and the introduction of new products.

Automated quality adjustment

As Figure 1 shows, automated indexes based on simple average prices across the data can easily lead to big problems. With a $2,500 bottle of wine appearing in the web scraped data, simply taking the average price causes a big discontinuity, when in reality this price change corresponds to a

[1] Multilateral methods utilise the data (prices and quantities, and in some cases also product characteristics) of products for all time periods within an estimation window, generally at least a year long. More traditional 'bilateral' indexes (e.g. the Laspeyres, Paasche, Fisher, Tornqvist) use prices and (sometimes) quantities of two periods - a base period and the current reference period.

[2] The fixed-effects aspect of this index is also commonly referred to within the Prices literature as the time-product dummy (TPD) index

significant quality change and should therefore be factored out as 'quality change'. A multilateral method (the FEWS in this case) easily controls for this in an automated way, without requiring that the expensive wine be identified as an outlier and removed or quality-adjusted using subjective approaches.

Figure 1. The impact of quality adjustment on price indexes

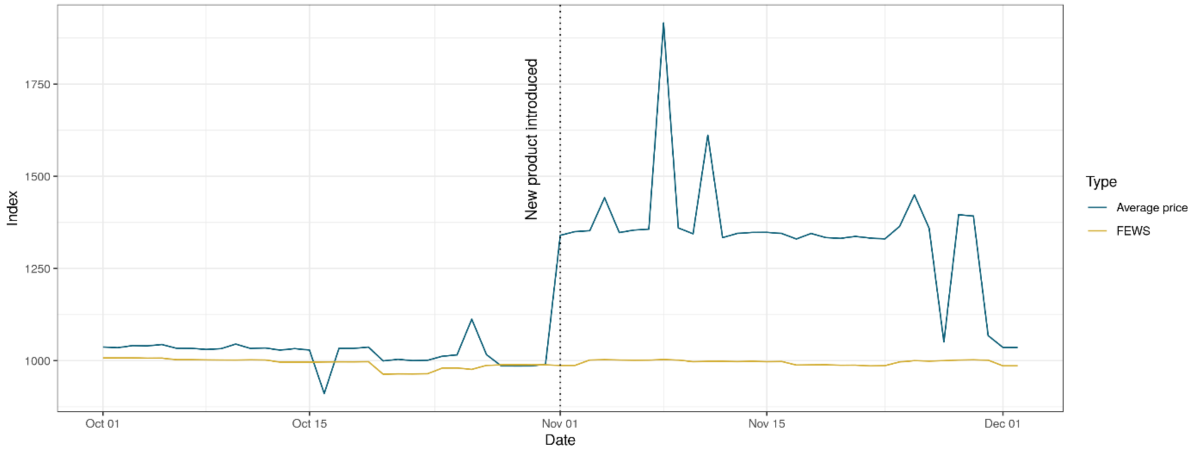

Product churn

Figure 2 illustrates the rapid churn in the set of products being bought and sold by showing, for seven different product classes, the percentage of the product types at the start of the period (July 2008) which are still being sold in each month over the following three years. Clearly, a fixed basket based on all products transacted at the start of the period would quickly become unrepresentative.

Figure *2: Percentage of July 2008 product specifications still available from July 2008 to June 2011*

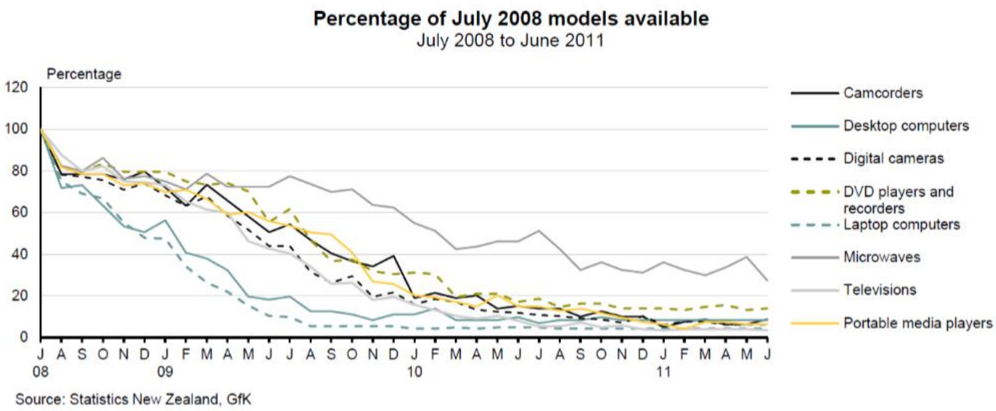

Chain drift

One solution to this dynamic product universe would seem to be the use of high-frequency chained superlative indexes – that is, using the prices and quantities in period t and period t-1 to estimate price movements with such methods as the Tornqvist index. These indexes could then be chained together. Unfortunately, such high-frequency indexes are affected by what's known as 'chain-drift', commonly held to be a consequence of the volatility of prices and quantities (due often to sales) not being perfectly symmetrical, although research is ongoing.

Figure 3, from de Haan & van der Grient (2009) shows the behaviour of actual (average) prices and quantities in Dutch scanner data for a particular type of tablet. This is the type of behaviour that is typically associated with chain drift in high-frequency chained superlative indexes.

Figure 3: Asymmetric spiking of prices and quantities in high-frequency data

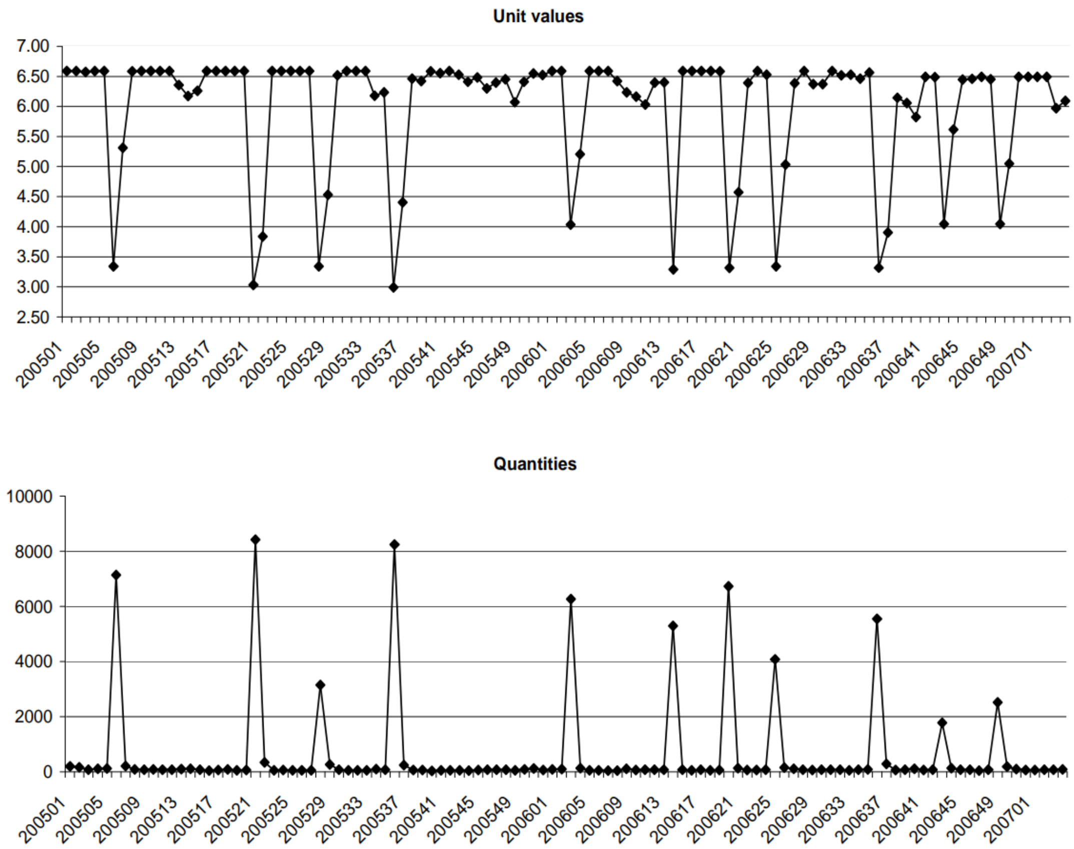

This chain drift is one of the reasons why new methods are required for alternative data. As mentioned above, there are a variety of multilateral index methods which have been developed and refined over the last few decades, which appear to perform well. Some of the methods require full information on characteristics to perform 'hedonic' quality adjustments, and some require only that prices and quantities are available for all products over time to enable product matching longitudinally.

Seasonality

One aspect of price data that the multilateral methods can cope with particularly well, compared to the traditional fixed basket approach, is seasonal products, where prices and quantities change markedly within a year. An example is fruit and vegetables, for which traditional methods have never been particularly adequate. Turvey (1979) illustrated the wide range of results when the same dataset was sent to different statistical agencies around the world, who were asked to apply their usual procedures.

For example, see figure 4, which shows the variation in monthly average prices, quantities sold and total expenditure across three years, for digital cameras, as measured from scanner data. Large spikes in quantities sold around December correspond to lower average prices, presumably because more lower priced cameras are being bought during the Christmas and Boxing day sale period.

Figure 4: Seasonality of quantity, expenditure and average price for digital cameras

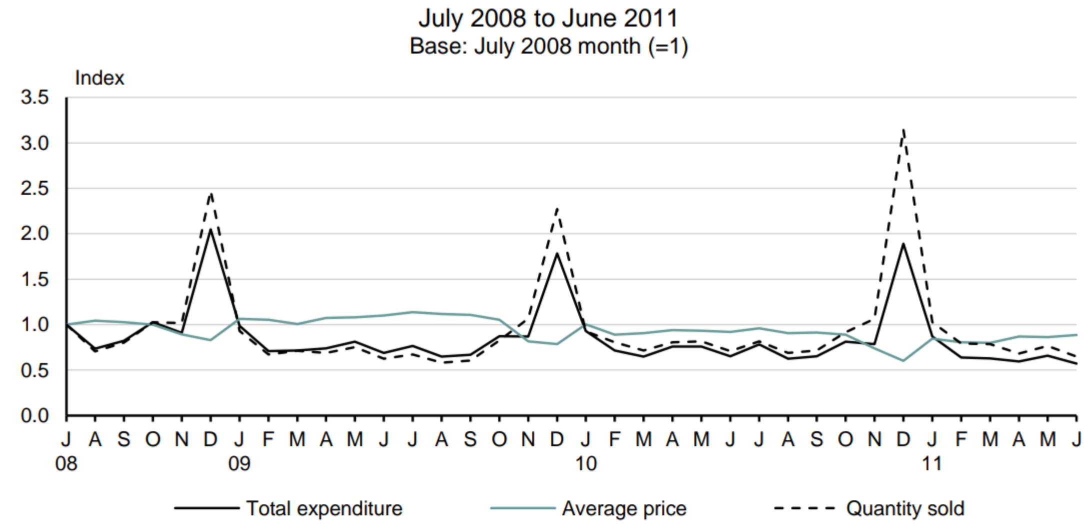

Note: This graph is based on earlier research data (Statistics New Zealand, 2012). Note that each of the three series is expressed as an index and based to 1 in July 2008, for confidentiality reasons.

Source: Statistics New Zealand

Figure 5 shows the result of the ITRYGEKS multilateral index, which automatically applies quality adjustment using 'hedonic regression' based on the full set of product characteristics contained in the scanner data. Note that after this quality adjustment the pure price change is no longer seasonal for this product.

Figure 5: Digital cameras - quality adjustment of price change with a multilateral method

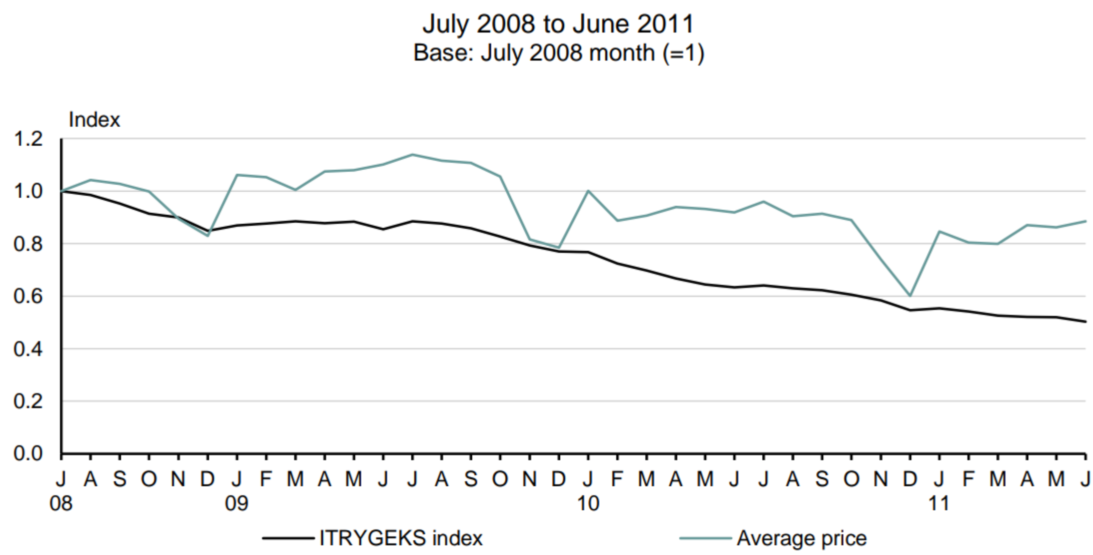

Note: It is possible, and valid, for quality-adjusted price indexes for some goods (eg fresh fruit and vegetables) to still display seasonality even after appropriate incorporation of seasonal quantities.

Source: Statistics New Zealand

Potential new alternative data sources

Expenditure data for supermarket products

Stats NZ currently receives average monthly prices for from the main supermarket retailers, for all products. This data enables substitution of scanner data for the products selected in the traditional fixed basket. Quantities sold (or, equivalently, total expenditures) would enable full use of supermarket scanner data using multilateral methods. Case study 1 below shows that expenditure weights can make a large difference to the price indexes, so Stats NZ would not consider full implementation of supermarket scanner data without these.

Web scraped online data

Since 2016 Stats NZ has been purchasing daily web-scraped online price data from PriceStats, the commercial counterpart of MIT's Billion Prices Project. This data captures, in near-real-time, online prices from a wide range of different NZ retailers.

With many other statistical agencies incorporating web-scraped online data into their price indexes, though generally using a 'substitution' approach rather than full-coverage use, Stats NZ wanted to get a good back-series of data with which to compare performance of price indexes from online data to traditionally collected indexes. This would enable hitting the ground running if and when Stats NZ decides to use online data, with a back series already available.

A big limitation of the web-scraped online data for full-coverage use is the lack of quantities. However, there is potential to use the data for real-time indicative inflation indicators, similar to the approach used by PriceStats. If Stats NZ was to do this, we would want to research whether to use a similar approach to that of PriceStats, who use a chained Jevons index, or whether to adopt a multilateral method such as the (unweighted) FEWS index which would allow for easier decomposition of price indexes to show contribution of different products to higher level indexes.

With a range of open-source web-scraping applications available, Stats NZ might decide to take a more targeted approach to web-scraping for substitution into the traditional fixed basket, and incorporate some analyst judgement into whether and how to scrape additional information on product characteristics to assist with quality adjustment when product replacements are required.

District Valuation Roll data for an official House Price Index

Stats NZ doesn't produce an 'official' New Zealand House Price Index. There are currently a number of different sources of house price indexes for New Zealand (REINZ, CoreLogic, QV) which all use a method called the Sales Price Appraisal Ratio (SPAR) index.

Direct use of consolidated district valuation roll (DVR) data on property valuations and sales would enable Stats NZ to use a multilateral approach (the FEWS) to produce house price indexes at a disaggregated regional level with the ability to identify contributions of different types of properties to the price movements. Previous research (Krsinich, 2019) showed the potential of this method for decomposition of the house price index into land and structure components.

Partnerships for in-house price indexes from confidential data

In some cases, rich sources of price data could be held by third parties who are unable to provide the detailed data to Stats NZ due to confidentiality restrictions.

In these cases, there may be potential to embed price index estimation processes within the third parties' IT environment in such a way that the price indexes can be calculated without confidentiality risks. Alongside price index estimation processes, it would be necessary to have monitoring and diagnostic processes running to ensure that any quality issues with the data were identified, or unusual price movements explained, without risk to the confidentiality of the underlying detailed data.

Generalising processing systems and methods

Currently, the processing of alternative data sources with non-traditional methods sits across different languages and operational systems, for example Excel, SAS, R and other bespoke systems.

Early processes, such as the index estimations for used cars and consumer electronics, are inefficient in various ways by today's standards. For example, the splicing[3] of the most recently estimated quarter's movement for used cars is done in Excel in a relatively manual way rather than coded into the production of the index as with later iterations of the methods for other products.

Also, the theory and practice of splicing has evolved, from the initial method of using the most recent movement only, to a variety of different approaches. These new methods can enable incorporation of updates due to newer data in the back period of the estimation window. Examples are the 'window splice' and the 'geo-mean splice' which incorporate revisions to the back periods of the newly estimated index.

By centralising the process and enforcing 'steady states' integration of new data sources and methods will be more streamlined. Creating a centralised system also brings transparency to these complex processes and a platform on which team members can learn, and where documentation and instructions can be readily linked.

While making production processes for existing use of alternative data more robust and transparent, this generalised system will also do much of the groundwork for future implementations of new data sources.

Figure 6. The proposed Multilateral Application Pipeline (MAP)

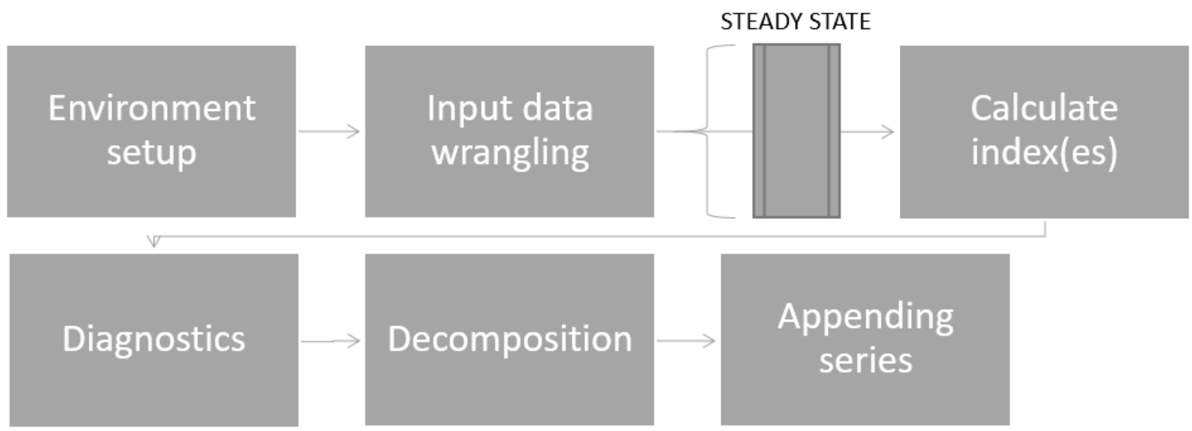

[3] The CPI is non revisable. Multilateral methods, however, re-estimate a back series with each successive period. This means that new results have to be 'spliced' onto the published index in such a way as to preserve the integrity of the published series (by incorporating a 'revision' factor). See de Haan (2015) for a discussion of different splicing approaches.

Empirical work to test methods

Recent work was done to determine whether and what can be generalised, additionally which methods are appropriate for different product types and data that might be utilised in the future. Some results from the work follow, but the general conclusions arising from the work are that:

- Expenditure weights are important, so we should prioritise getting expenditure weights to enable full-coverage use of supermarket data.
- The importance of expenditure weights also means that we shouldn't make use of full-coverage (unweighted) web-scraped online data, but rather use it for substitution into a fixed basket and/or use for indicative real-time inflation indicators.
- Of the weighted methods, FEWS and GEK-T generally give very similar results to each other, and to the benchmark ITRYGEKS method, for the products tested.
- When no weights are available, FEWS appears to give less volatile results than the alternative methods (GEKS-T and GEKS-J), therefore we would be likely to use this method for any indicative real-time indicators from online data.

Case study 1 – scanner data for supermarket products

For an examination of different methods' performance on weighted supermarket scanner data we used six years of weekly aggregated US supermarket scanner data from IRI marketing (see Bronnenberg et al. 2008). This data includes prices and quantities sold but no information on product characteristics. For manageability, and to reduce the unnecessary volatility of a weekly index, the data was pre-aggregated to a monthly level for the analysis.

Note that without a full set of price-determining product characteristics in the data, explicit hedonic indexes such as the time-dummy hedonic (TDH) index or the Imputation Tornqvist Rolling Year GEKS (ITRYGEKS) index are not possible.

Instead, for each product we produce the FEWS and GEKS multilateral methods. We base the GEKS on two different underlying bilateral calculation methods, the Tornqvist (GEKS-T) and the Jevons (GEKS-J).

The GEKS-Jevons index is, by definition, unweighted, but for both the FEWS and GEKS-Tornqvist indexes, we produced indexes with and without weights.

Weights are important for supermarket scanner data

For some supermarket products we found the removal of weights had a significant effect, for example blades[4], shown in figure 7.

With weights, the GEKS-T and the FEWS give almost identical results.

Without weights, the trends of the indexes are significantly flatter and there is divergence across the three methods, with the GEKS-T significantly steeper than the GEKS-J and the FEWS.

[4] The 'Blades' product group doesn't include razors, which are a separate product in the IRI data, so it presumably consists of products such as knives and scissors.

Figure 7: Blades – monthly price indexes

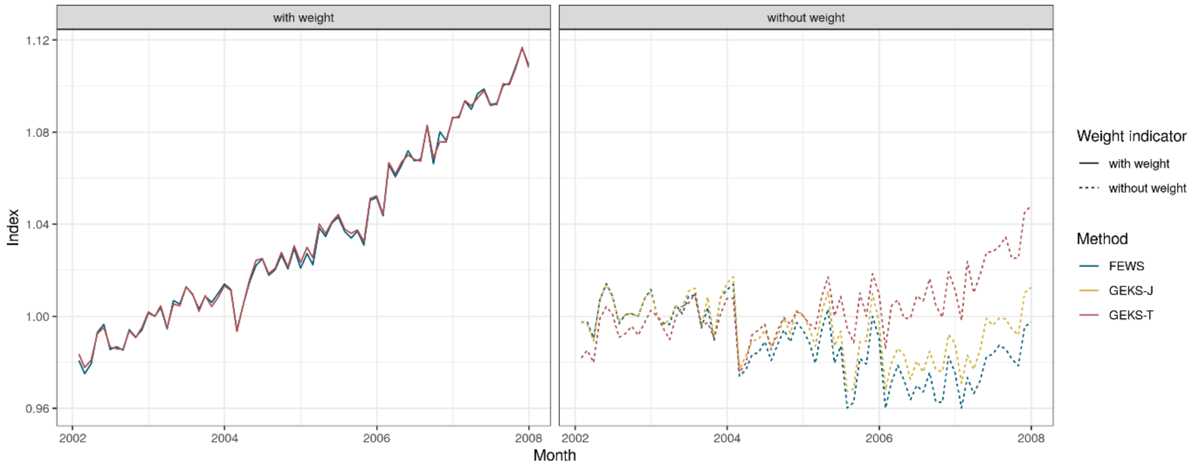

In contrast, some products such as Beer (shown in figure 8) show very similar long-term trends whether or not weights are incorporated. The weighted indexes are more volatile than the unweighted indexes.

Figure 8: Beer – monthly price indexes

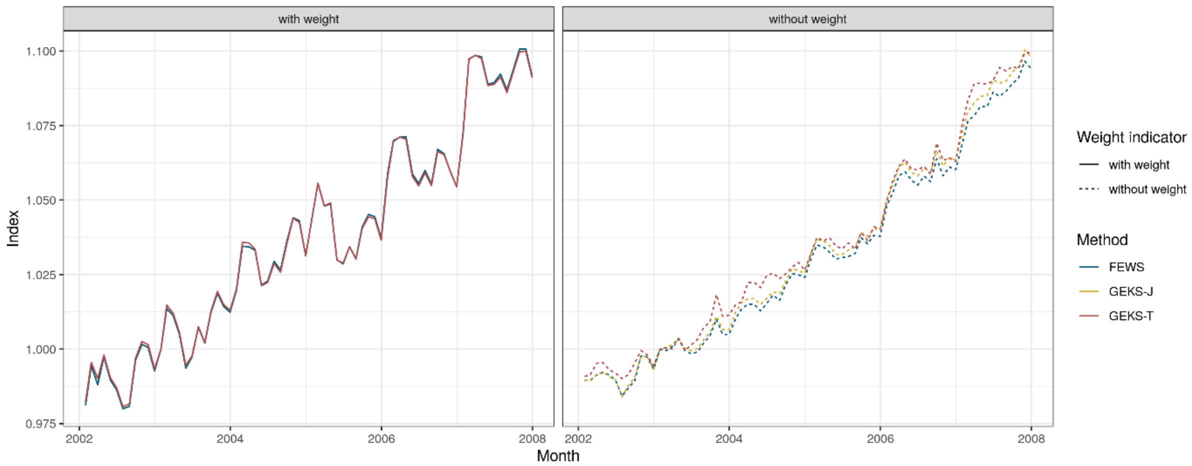

Similarly, figure 9 shows that weighted indexes for salty snacks have the same trend as, but are more volatile than, unweighted indexes. Graphs are grouped by method to highlight this pattern.

Figure 9: Salty snacks – monthly price indexes

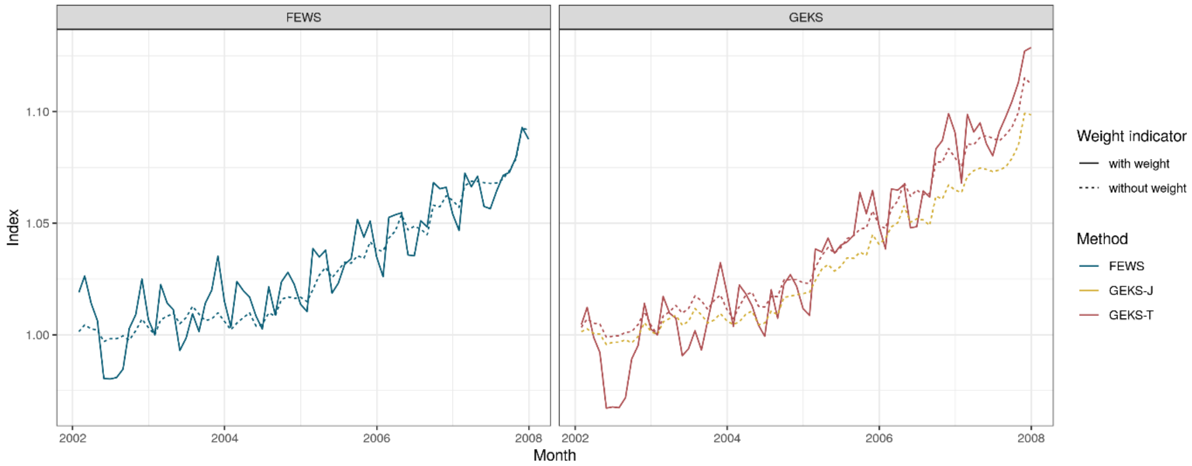

Figure 10 shows the comparison for shampoo, for which the weighted and unweighted methods are similar for the first three years and then diverge for both the FEWS and GEKS methods. Unweighted GEKS-Tornqvist and GEKS-Jevons behave quite differently, trending in opposite directions in the 2nd half of the six-year window.

Figure 10: Shampoo – monthly price indexes

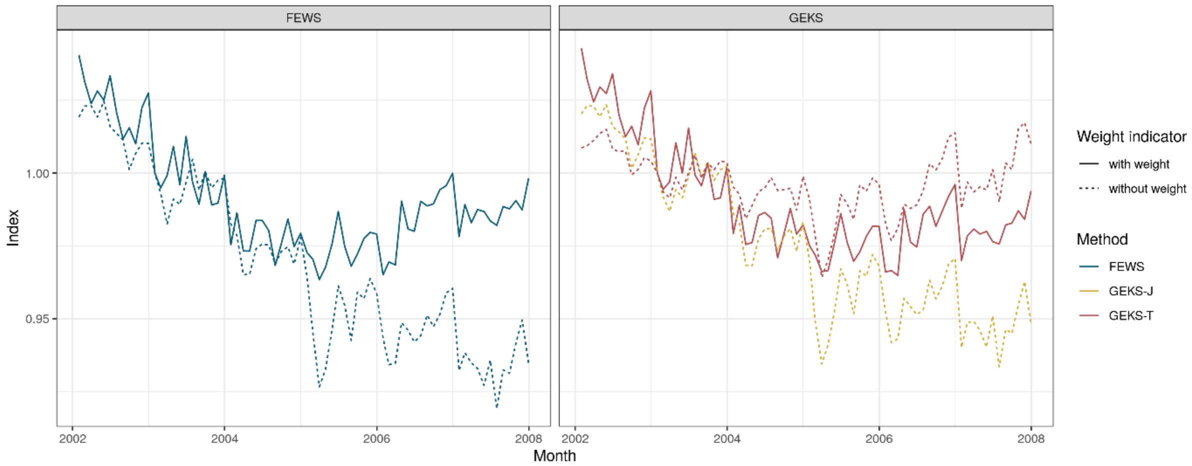

More work is needed to understand whether we can characterise the products who are affected by weights. It may be due to the elasticity of substitution. For example, beer seems likely to be less elastic than blades, with more potential for brand loyalty – for these products with less substitution happening it seems plausible that the weights across specific products will be less variable and therefore unweighted indexes will behave similarly to weighted indexes.

If we can't obtain expenditure weights for supermarket scanner data, understanding when weights are important becomes a higher priority as unweighted scanner or online data may be the only source of alternative price data.

Weighted GEKS-Tornqvist and FEWS indexes are highly correlated

Conceptually, we have tended to favour the FEWS multilateral index as it is based on a statistical model and appears to be appropriately adjusting for the introduction of new products when used with a splicing method such as the window splice or geo-mean splice which incorporates revisioning.

We note that the Australian Bureau of Statistics has decided to use a GEKS-Tornqvist for their implementation of supermarket scanner data, after a lot of investigation into the different options.

Our analysis of 31 IRI supermarket products shows that the weighted FEWS and GEKS-T give very similar results (see figure 11), so at this stage it seems likely that either method would be as good as the other. The generalised system we will develop will make testing and changing methods straightforward so, given that research into the details of the different methods' behaviour is ongoing, we can have the option to revisit and change decisions further down the track.

Figure 11: Correlation to monthly and annual change of weighted FEWS index

Each point in the scatterplot represents a different method applied to a different product. The x-axis is the Pearson correlation coefficient between monthly changes in the method of interest and the weighted FEWS and the y-axis is the correlation of the annual changes with those of the weighted FEWS.

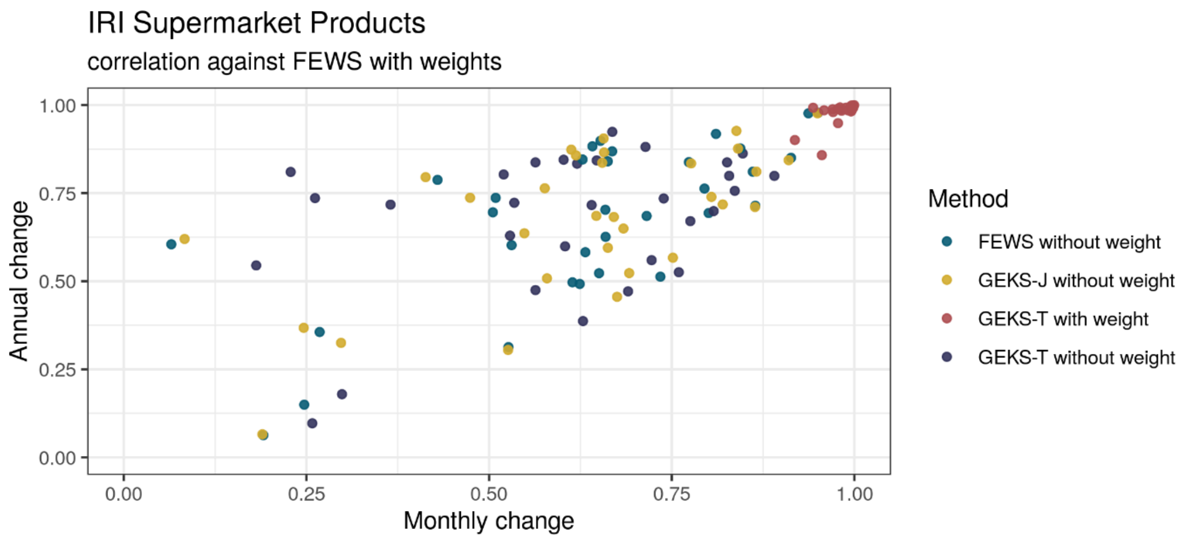

Case study 2 – scanner and online data for consumer electronics products

Price movements of consumer electronics products have been estimated in the NZ CPI using scanner data on prices, expenditures, and detailed characteristics since 2014.

A multilateral method called the 'Imputation Tornqvist Rolling Year GEKS' (ITRYGEKS) (de Haan and Krsinich, 2014) is used, which combines hedonic regression models incorporating product characteristics with a multilateral index method – the rolling year GEKS (Ivancic et al, 2011). At the time of its introduction this was considered the 'gold standard' method given its explicit hedonic quality adjustment using product characteristics and the drift-free nature of the GEKS aspect of the index.

Weighted multilateral methods without characteristics give similar results to the ITRYGEKS benchmark

Using the ITRYGEKS as the benchmark, we can compare the performance of other, simpler, multilateral index methods which use only prices and quantities at the detailed product specification level. That is, against the ITRYGEKS benchmark we compare methods which don't explicitly incorporate product characteristic information.

We have scanner data (prices and expenditures) for consumer electronics products from market research company GfK, and web-scraped online data (prices but no quantities or expenditures) for some of the same products from PriceStats. Note that the coverage of the PriceStats data is different as not all retailers are included, however we would expect to see similar price movements.

This gives us the opportunity to test how well online data would work for these products, which are at the most challenging end of products in terms of having:

- High product churn
- Short life cycles
- Potential for life-cycle pricing, in terms of both 'new product premiums' and end-of-life discounting
- Rapid technological innovation across time between products

Figure 12: Desktop computers – monthly indexes

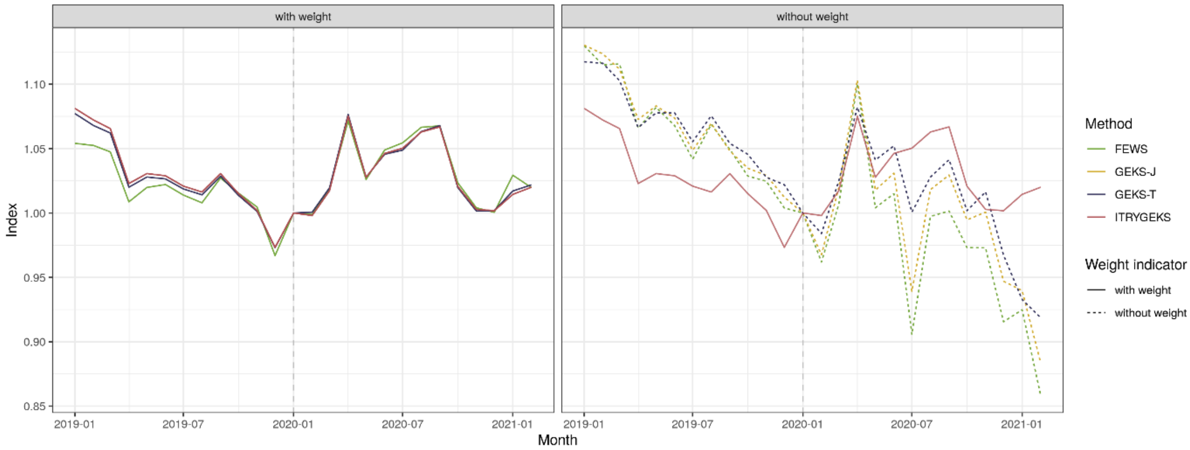

Figure 12 compares the FEWS and GEKS-T to the ITRYGEKS, on the same scanner data from GfK. It shows that, despite not having any information on product characteristics with which to do explicit quality adjustments, the (weighted)FEWS and GEKS-T give very similar results as the ITRYGEKS.

When expenditure weights aren't included, however, the unweighted FEWS, GEKS-T and GEKS-J (which is unweighted by definition) can significantly differ, from both the ITRYGEKS and each other.

This demonstrates what we found above for the supermarket products using IRI scanner data – when weights are included, there is not much difference across the ITRYGEKS, FEWS and RYGEKS-T indexes. Our conclusion is again that expenditure weight information is important.

Figure 13 demonstrates the same point – that expenditure weights are important – by showing correlations with ITRYGEKS indexes for both monthly and annual change, across products and methods. The weighted GEKS-T corresponds most closely to the ITRYGEKS, followed by the weighted FEWS.

Figure 13: Correlation to monthly and annual change of ITRYGEKS index

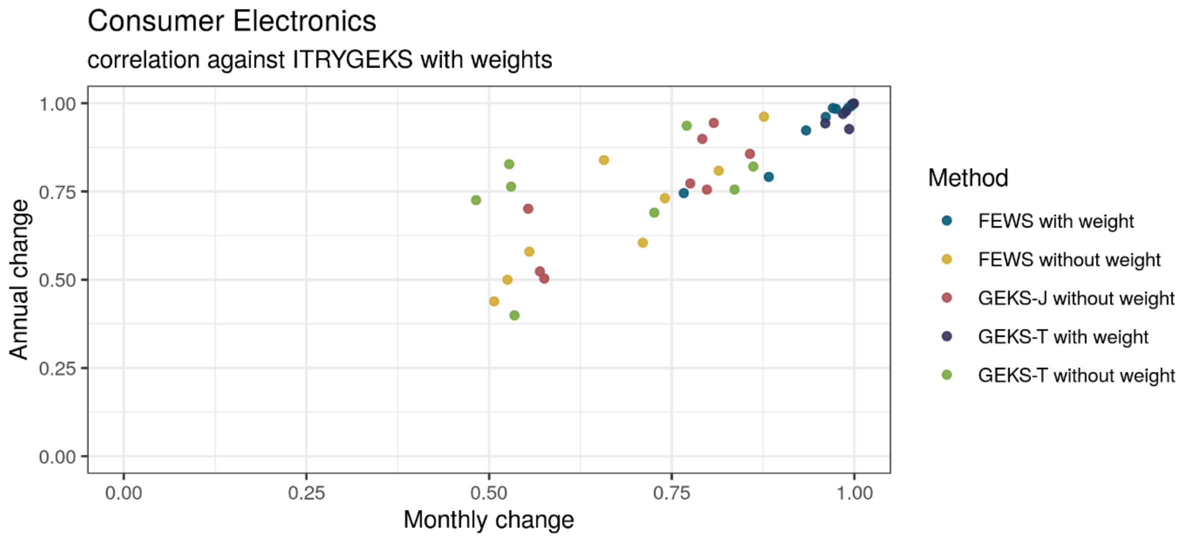

One of the potential data sources we're looking at is web-scraped online data. This will generally have lower coverage than the scanner data, as not all retailers have an online presence. We won't have the corresponding expenditure weights, and there is some potential for online prices to differ from in-store prices.

So we would expect divergence in unweighted indexes from online data, against the corresponding weighted scanner data for the same products. But we can empirically test this for the products which we have from both sources.

To summarise, the difference between indexes from online data and the benchmark index ITRYGEKS from scanner data and the other methods could be due to any or all of:
- the method
- the weights
- the coverage
- how similar online prices are to instore prices.

First, looking at the effect of just the different methods, we compare FEWS, GEKS-J and GEKS-T for four different consumer electronics products – desktops, headphones, laptops and televisions.

The methods all give very similar results for headphones, but for TVs there are significant differences between the unweighted FEWS and the unweighted GEKS-T and GEKS-J, which are significantly more volatile. If using unweighted data, we'd want to be more confident about which of the GEKS or FEWS indexes was conceptually more appropriate, and this would require further work.

Figure 14: Online data with prices but no quantities sold – four consumer electronics products

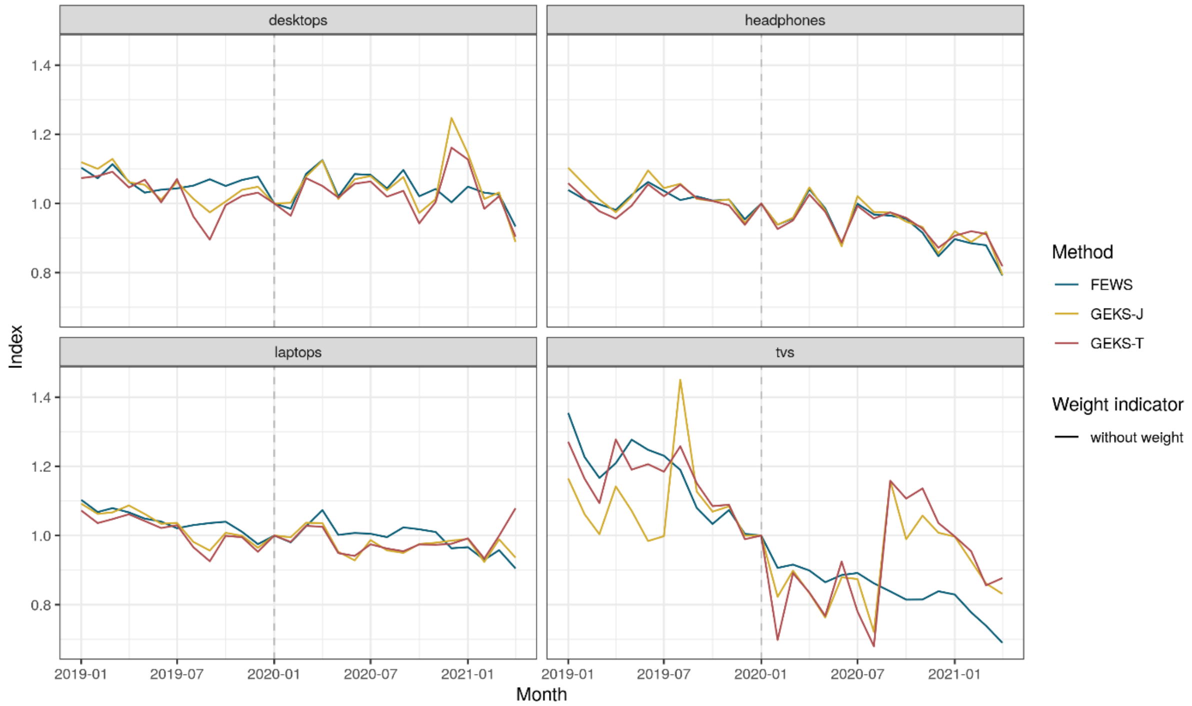

Figures 12 and 14 above demonstrated that when weights are included there is very little difference between the benchmark ITRYGEKS index and each of the FEWS and the RYGEKS.

Figure 15 below shows the aggregate effect of the weights, the coverage and how well online prices represent instore purchases, by comparing the weighted benchmark ITRYGEKS index to unweighted indexes from online data for desktop computers. For this product, at least, the FEWS index is sitting closer to the benchmark index, so might be more appropriate than the GEKS-T or GEKS-J.

Figure 15: Desktop computers - comparing unweighted indexes for online data to weighted scanner data benchmark index

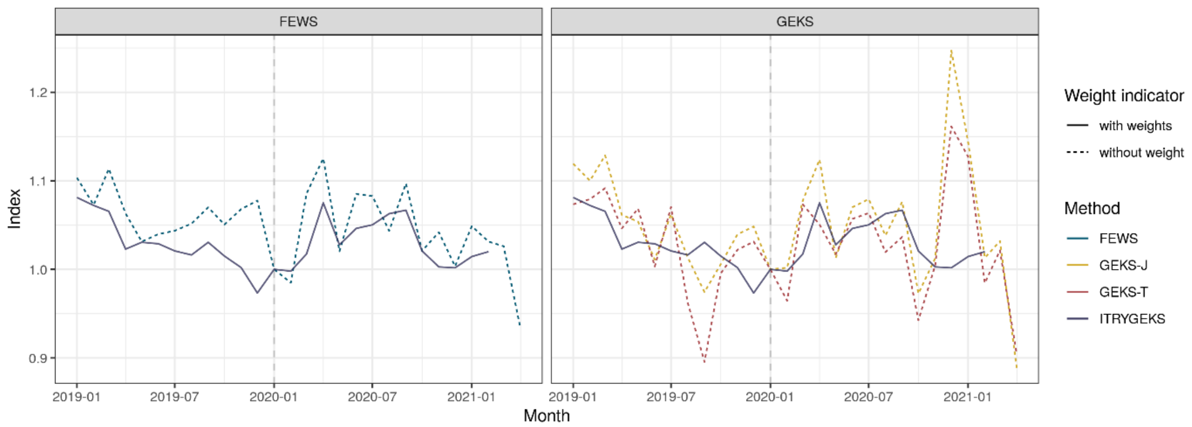

Lack of weighting means we would be unlikely to use online data with multilateral methods for full-coverage incorporation into the NZ CPI. It’s use is likely to be via substitution into the traditional basket. However, for indicative real-time inflation indicators similar to PriceStats’ measures, this research suggests that the FEWS may be the appropriate index.

Aggregation of indexes

When multilateral indexes are introduced into the CPI, we need to consider at what level they are introduced and how they link with any traditionally collected prices (e.g. field) we are still using.

The easiest method is to implement them is at the lowest level of elementary aggregation in the CPI structure for each store as this would require no changes to the way we structure our indexes. In Stats NZ CPI this is a regional index for a specific item in the CPI basket (e.g. Oranges in Wellington).

However, there are other options that should be considered when introducing multilateral indexes:

- Introducing at a higher level in the CPI aggregation (e.g. class level – fruit)
- Aggregating multiple stores together when calculating the multilateral movement
- Aggregating all regions together to calculate a NZ movement.

Each of the above options has pros and cons. For example if we introduce multilaterals at a higher level then we account for seasonal items by having dynamic weights, dampening the effect of volumes decreasing when prices increase which gives a more accurate movement for that class. The downside to this is that we would need to create the lower level indexes separately as these are needed for specific customers and used in other areas of Stats NZ.

Conceptually, our preference is to introduce the multilateral index at the class level (where possible), aggregating multiple stores together in each region. This would account for substitution between products and stores within a class and region. There would not be much substitution between regions and having regional indexes is an important measure, so keeping a regional split is preferred. An example of how this would look for fruit:

Figure 16. The structure of the index for fruit

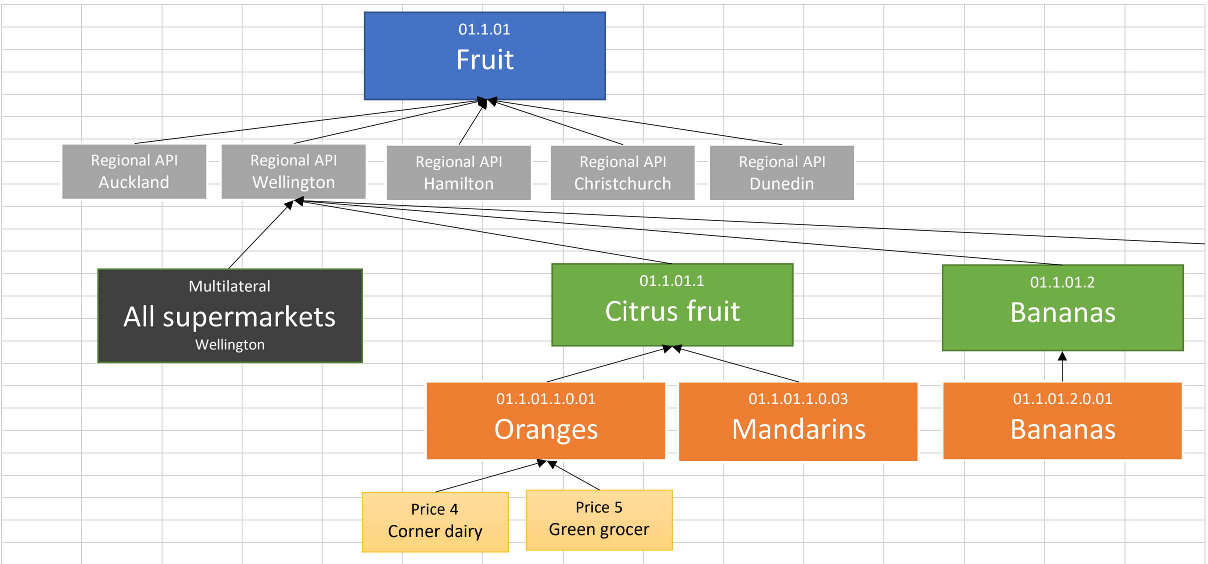

Figure 16 shows how these multilateral indexes might link with traditionally collected prices. Because traditionally collected prices cannot be used in a multilateral index, they would need to sit beside the multilateral indexes and feed into a more traditional CPI structure.

We will be doing more investigations on the above options in the coming months, aiming to confirm our preferred approach and quantifying the impact this change would have on our release products.

Conclusion

The use of alternative data for price indexes, in terms of both methodology and processes, is still an active area of research both here in New Zealand and internationally.

There as increasingly well-established body of research and development on multilateral methods for use with the kinds of full-coverage big data that we need to deal with when utilising scanner data, online data or administrative data for price measurement.

The development of a generalised system for estimation and analysis of price indexes from alternative data is our best use of resource. It will make existing production processes more robust, provide a good platform for documentation and staff training, and set in place 80% of production processes for any future use of new alternative price data sources.

References

ABS (2017) An implementation plan to maximise the use of transactions data in the CPI. Information Paper. Australian Bureau of Statistics

Bentley, A (2018) Rentals for housing: A model-based estimator of inflation from administrative data Paper presented at the Meeting of the Group of Experts on Consumer Price Indices, Geneva, Switzerland

Bentley, A & F Krsinich (2017) Towards a big data CPI for New Zealand Paper presented at the Ottawa Group 2017. Eltville, Germany

Bronnenberg, B.J., M. Kruger, and C.F. Mela. 2008. The IRI Marketing Data Set Marketing Science 27: 745 – 748. (accessed April, 2016).

de Haan, J (2015) Rolling Year Time Dummy Indexes and the Choice of Splicing Method. Paper prsent

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. Earlier version available here

de Haan, J & H van der Grient (2009) Eliminating Chain Drift in Price Indexes Based on Scanner Data

Ivancic, L, W.E. Diewert and K.J. Fox (2011), "Scanner Data, Time Aggregation and the Construction of Price Indexes", Journal of Econometrics 161, 24-35.

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 conference, Wellington, New Zealand

Krsinich, F (2014) Quality Adjustment in the New Zealand Consumers Price Index. Chapter from 'The New Zealand CPI at 100. History and Interpretation'. Publisher: Victoria University Press. Editors: Sharleen Forbes, Antong Victorio

Krsinich, F (2016) The FEWS index: Fixed effects with a window splice. Journal of Official Statistics. Vol 32, issue 2

Krsinich, F (2019) Land prices: UNCOVERED! Extricating land price indexes from improved property price indexes for New Zealand. Paper presented at NZAE conference. Victoria University, Wellington.

Stansfield, M (2019) Import and export price indexes using fixed-effects window-splicing NZAE conference

Stats NZ (2014) Measuring price change for consumer electronics using scanner data

Stats NZ (2019a) New methodology for rental prices in the CPI

Stats NZ (2019b) Overseas trade price indexes through a multilateral method

Turvey, R. (1979), "The Treatment of Seasonal Items in Consumer Price Indices," Bulletin of Labour Statistics, ILO

UNESCAP (2019) Incorporating non-traditional data sources into official statistics – the case of consumer price indexes. Lessons and experiences from Australian, Japan and New Zealand