

Creating a digital food price index from web-scraped data

Paper presented at New Zealand Association of Economists,
at Auckland, New Zealand,
27 June 2018

Donal Lynch

Matthew Stansfield

Sam Olivecrona

Statistical Analysts, Prices, Accommodation, and Construction, Stats NZ
P O Box 2922
Wellington, New Zealand
info@stats.govt.nz
www.stats.govt.nz

Disclaimer

Conference papers represent the views of the authors, and do not imply commitment by Stats NZ to adopt any findings, methodologies, or recommendations. Any data analysis was carried out under the security and confidentiality provisions of the Statistics Act 1975.

Liability statement

Stats NZ gives no warranty that the information or data supplied in this paper is error free. All care and diligence has been used, however, in processing, analysing, and extracting information. Statistics New Zealand will not be liable for any loss or damage suffered by customers consequent upon the use directly, or indirectly, of the information in this paper.

Reproduction of material

Any table or other material published in this paper may be reproduced and published without further licence, provided that it does not purport to be published under government authority and that acknowledgement is made of this source.

Citation

Lynch, D, Stansfield, M, & Olivecrona, S, (2018, June). *Creating a digital food price index from web-scraped data*. Paper presented at New Zealand Association of Economists, Auckland, New Zealand.

Abstract

How do you turn web-scraped data into something useful? As part of Stats NZ's initiative to explore alternative methods of collecting data for price indexes we are investigating web scraping supermarket prices and transforming them into a real-time 'Digital Food Price Index'. Stats NZ recently up to purchased daily web-scraped online price data from PriceStats, the commercial counterpart of Massachusetts Institute of Technology's (MIT) Billion Prices Project. This data captures online prices from a wide range of different NZ retailers. Preliminary research on measuring food price change from this data is used to illustrate our approach to using 'big data'. This paper gives an overview of Stats NZ's progress into web scraping for price collection and highlights the opportunities and obstacles of coverage, timing, and quality adjustment associated with using 'big data'.

1 Introduction

Many national statistical agencies have gained access to retailer transaction data. This data consists of a time series of all items purchased in store and the amount paid by the customer. One key feature of this data is that it contains the volume of sales in addition to the prices. This data is commonly referred to as scanner data. According to Chessa et al. (2017) there are currently 7 European countries incorporating scanner data into their published Consumer Price Index (CPI), as well as Australia. Many other countries have access to the data and are actively developing methods to implement the data into their CPI calculations.

In addition to scanner data, web scraping has been brought to the attention of government statistical agencies for some time. The rise of online shopping has led to a rise in the potential for price collection via web scraping. As such, research in the area of online price collection and analysis through automated scrapers is very active. This is also a frequent topic among the UN Ottawa group members (international experts in price indices), appearing either as talks within big data sessions or as an entire session on its own. Additionally, a new section on scanner data is being added to the International Labour Organisation CPI manual.

As noted above, the primary difference between scanner data and web-scraped data is that scanner data has prices and quantities of purchases, whereas web-scraped data only contains the price information. Although Stats NZ does have scanner data from some retailers, the agency does not have any scanner data from supermarkets or for any food prices. Progress has been made in implementing some scanner and administrative data into published indexes. Krsinich (2015) uses the ITRYGEKS¹ model to calculate consumer electronic price indexes and Krsinich (2011) uses administrative data and a THD² model to produce an index for used cars. Bentley and Krsinich (2017) present case studies on all current applications of "big data" in published Stats NZ statistics. As an extension of this, Stats NZ has been investigating the viability of implementing web-scraped data into the compilation of the food price index (FPI).

¹ Imputation Törnqvist rolling year GEKS

² Time Hedonic Dummy. More information on both of these methods in [section 2](#).

2 Background

Stats NZ uses a fixed basket of goods in our current production method for the FPI. This basket currently contains 162 items. Field officers visit supermarkets to collect prices for these items. Fresh fruit and vegetable prices are collected weekly, otherwise prices are collected monthly. If an item is not available, prices are imputed from the previous period. If the item is not available for a second consecutive period, the item is replaced with one of equivalent quality. This act of substituting an item with an item of equivalent quality and size is an important feature of the FPI. The FPI aims to measure the change in food prices while keeping quality and quantity constant.

The reality is that the replacement item is chosen based on the officer's best judgement, and so there is a subjective element to the replacement. The collected prices are then checked before being passed through a multiple stage quality assurance process. Because the sample of items is relatively small and the frequency of collection is relatively low, this data is easy to maintain and post-collection tasks are manageable for analysts.

The current approach is generally referred to as a bilateral method. A bilateral method involves comparing all prices from one period back to a single reference time period (the base period). One drawback with a bilateral method is that they cannot easily adapt to the introduction of new or removal of old items, hence the fixed basket of goods for the FPI.

Motivation for exploring web-scraped data

- Current field collection uses a small selection of products which are judged to be representative of all products purchased by consumers. However, consumers' preferences and needs constantly change, and the fixed basket concept cannot easily accommodate this. The contents of the basket are periodically reviewed to account for changes in spending patterns.
- Emerging international best practice is to move away from a fixed basket.
- Using a richer data source, such as web-scraped data, would offer more ability to draw out the drivers of change in the FPI and generate new insights.
- Web-scraped prices offer increased efficiencies in producing an index, with fewer field collections and less post collection analysis.
- When a new data source or method is presented, Stats NZ should "adopt, adapt, or explain". In this case web scraped prices are in the public domain. If Stats NZ chooses not to use this data for a food price index they should have a good reason as to why they are not adopting this approach. Similarly, if any other group produced an index based on web-scraped data, Stats NZ should be able to explain any discrepancies between that and the published index.

The use of web-scraped data offers the opportunity for real-time updating and the potential to do away with the concept of a fixed basket. However, there are also challenges with fixing quality and dealing with the fact that we do not know the quantities sold. Bilateral methods are not well suited to web scraped data where products may frequently drop out, and new products come in.

Multilateral methods are the alternative to bilateral methods. Multilateral methods compare the current time period to recent time periods, and so accommodate new and disappearing products. Table 1 summarises the suitability of multilateral methods for calculating indexes based on two criteria:

- Are the characteristics available for each item? For example, in the case of used cars the available characteristics include features such as the age of the vehicle, the engine capacity, the number of kilometres on the odometer etc. Characteristics are not available for the web scraped supermarket data.
- Does the method control for new and disappearing prices? For the supermarket scanner data we would ideally use a method which does account for changing products.

Table 1

Comparison of multilateral methods			
	Characteristics required?	Reflects implicit price movements of new/disappearing products?	Examples of current New Zealand production uses
Time hedonic dummy (THD)	Yes	Yes	Used cars
Rolling-year GEKS (RYGEKS)	No	No	
Imputation Törnqvist RYGEKS (ITRYGEKS)	Yes	Yes	Consumer electronics scanner data
Fixed-effects window-splice (FEWS)	No	Yes	Overseas Trade Index (phones/TVs)

3 The FEWS method

Based on table 1, a FEWS method is the most appropriate method to apply to the supermarket web scraped data. Krsinich (2016) introduced the fixed effects window splice (FEWS) index method as an alternative to methods like RYGEKS, which can introduce bias when products disappear from a series, or new products are introduced. The FEWS method is an amalgamation of the fixed effects (FE) model and a splicing approach (window splice - WS) that takes into account the price observations from across a fixed time window.

Fixed effects

The FE index, also known as time product dummy (TPD) is suitable for products where limited or no characteristic information is known, but there is a unique identifier available for each item. It has been shown by Krsinich (2016) that the FE model is equivalent to a fully interacted time-dummy hedonic index based on all price-determining characteristic information, despite those characteristics not being observed.

In essence, the FE index is calculated on the change in the price of each item, compared over a given time window. In this way, controlling for quality is implicit as any change in an item, such as packaging, quantity or any other quality should result in a new unique identifier for that item.

For a product to contribute to the index it must have at least two observations in each window. However, due to the nature of scanner data or web scraped data, which has many observations of each item, this tends to be an easily satisfied constraint.

Splicing

In the TPD method a fixed window length is chosen. The choice of window length can be non-trivial. Bentley (2018) found that the choice of window length significantly influenced an index for home rental prices, whereas Chessa et al. (2017) found that window length choice had a small effect on scanner data. Our findings for the effect of window length on web-scraped data are shown in [section 5](#). An FE index is then calculated over that window. The window is shifted forward by one time period, and the index is recalculated. Thus there is an overlap between the windows of (*window length - 1*) time periods. In order to form a continuous time series index, these windows must be chained together, while maintaining the no revision constraint which applies to most national consumer price indexes.

Ivancic, Diewert, & Fox (2009) suggest splicing the windows together based on the final time period as the best approach. This became known as the movement splice. Krsinich (2016) suggests a better approach is to use the second time period in the window to splice the windows together, and calls this a window splice. De Haan (2015) suggests that the time position for splicing should be the middle of the window, which became known as the half splice. Diewert and Fox (2017) point out that an argument can be made for any time period in a window to be a suitable splice position, and the “best” approach is to take the geometric mean of all possible splice positions.

4 Web scraping and the PriceStats data

Stats NZ currently purchases web scraped price data from PriceStats, the commercial counterpart of MIT’s Billion Prices Project. The dataset consists of daily web-scraped data from 14 New Zealand online retailers, each of which also have physical stores. In our work developing a digital food price index we have focused on the data from two supermarkets. The supermarket data dates back to early 2014 for one of the retailers and mid 2016 for the other. The data is collated into daily tables with one row per price observation and the following columns, see table 2.

Reasons for focusing on the FPI

There are a number of reasons why food is a good choice for developing our ability to build a price index from web scraped data:

- The monthly FPI is of national interest, and with about 20% of weight is a leading indicator for the quarterly CPI.
- Potential to produce index more efficiently.
- The coverage of food prices by online retailers is quite good.
- Food products have lower churn rates than other items.

These bullet points will be expanded on in the following sections.

National Interest

Food is the only component of the CPI which is published on a monthly basis. The reasons for this are:

- A quarterly index does not carry enough resolution for the seasonality of food prices and their fluctuations.
- The price of food is important to many government agencies and New Zealand residents alike
- The FPI has a weight of about 20% of the CPI, and is a leading indicator for the quarterly CPI.

Table 2

Column details from PriceStats dataset	
Column name	Details
Retailer	The name of the store from which the data is sourced
Price date	The date of the web scraping
RID	A unique number to identify the product (e.g. SKU ¹)
Product name	The description as per the website
Price	The price of the item
Price no sales	The 'off sale' price of the item
Level 1 – Division	Price stats employ a machine learning algorithm to match each item, based on description, to the equivalent COICOP ² classification, coded to level 3.
Level 2 – Group	
Level 3 – Class	
Confidence	The confidence of the COICOP coding for each item
Product URL	The URL of the item
Retailer menu	Websites usually have a hierarchical structure to organise their products. For example, a customer looking to purchase bread might go through the following categories on a retailer website: Baked goods -> Sliced Bread-> White bread. This hierarchical trail is known as the retailer menu.
<p>1. SKU = Stock keeping unit. This code is unique to each item for sale. These codes are not regulated or standardised.</p> <p>2. COICOP = Classification of individual consumption by purpose. Developed by United Nations Statistics division, COICOP is a standard classification system for expenditure items. COICOP is a structured system with each level giving finer item classification.</p>	

Potential efficiencies

In using web scraped data to produce an FPI there would be fewer visits to stores and less post collection analysis required. As part of our examination of the data we have also considered the possibility of directly substituting web scraped data into the current FPI methodology. This alternative mode of collection could only be used in those stores with online shopping available. Field officers would still be required for stores with no online shopping option. For this option to be viable, the following conditions must be met:

- supermarket prices do not vary significantly by region
- online prices are the same as instore prices.

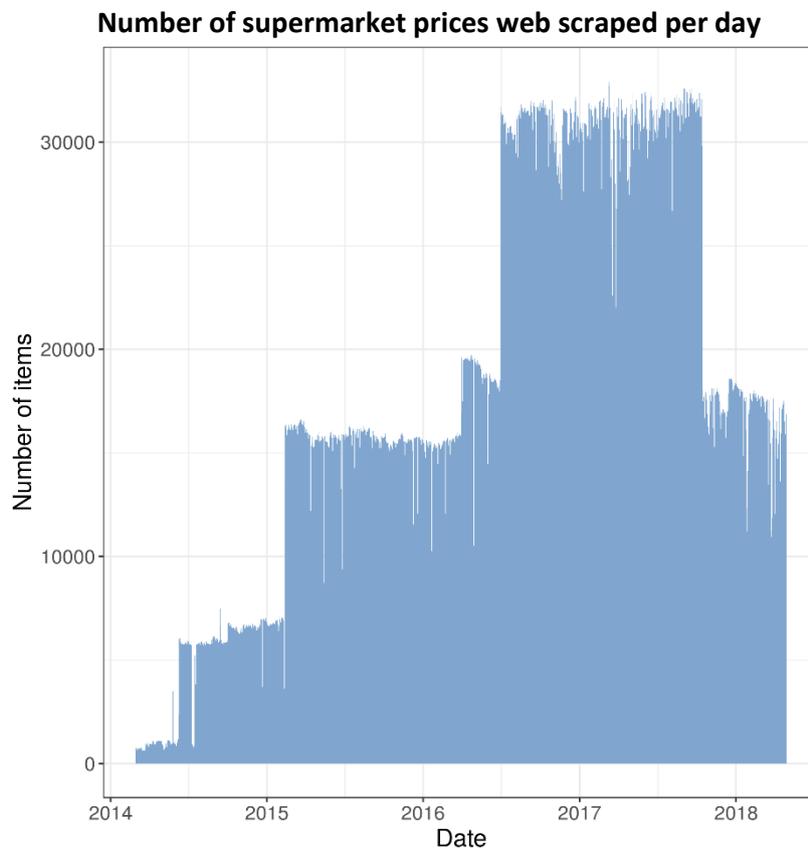
Research carried out in Australia by Waschka, Milne et al. (2003) shows that online supermarket prices have very little variation; neither regionally nor between online stores and physical locations. Cavallo (2017) conducted a large-scale review of online prices versus in store prices in 10 countries. He found that “Price levels are identical 72% of the time” and that price changes are not synchronised but have similar frequencies and average amounts. At Stats NZ we have also undergone an assessment of the price differences between field-collected supermarket prices and web scraped prices. Our findings match those of the Australian Bureau of Statistics (ABS). Most price discrepancies between online and field collected prices can be accounted for by the uncertainty in the date of field collection. Sally Copland, head of online at Countdown, stated that prices are the same online as in store (Gibson, 2017).

Good online Coverage

The prices of up to 30,000 New Zealand supermarket items are available online each day, and scraped by PriceStats. Of these, 43% are in the “Food and non-alcoholic beverages” division. Figure 1 shows the number of daily prices gathered from New Zealand supermarkets by PriceStats. Although a large number of observations are made per day, the noise in the number of prices is reflective of some of the inherent issues with web scraped data.

Figure 1 shows some of the inherent risks of using web scraped data. As a national statistics agency we must be confident of data being available for published indexes. Whether this threshold of confidence can be met by supermarket website data remains to be seen.

Figure 1



There are several step changes in the graph. The details of these are outlined below:

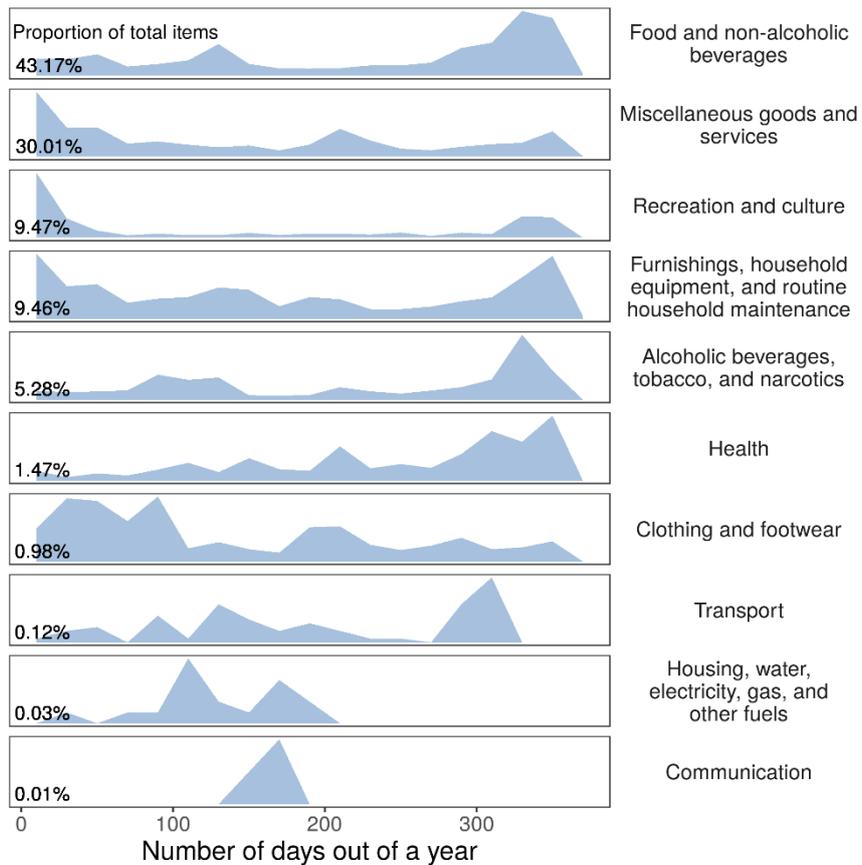
- 2014 Q3: PriceStats refining their web scraping method. Data only collected from one retailer
- 2015 Q1: Further improvements to the PriceStats web scraping routine
- 2016 Q1: More prices introduced by one retailer
- 2016 Q2: A second supermarket makes prices available online
- 2017 Q3: One retailer removes their website sales and moves to app based system. Neither Stats NZ nor PriceStats currently have the capability to scrape apps.

In addition to the step changes due to observable events, there is a large day to day variability. Some potential reasons for discontinuous data may be:

- The website changing its layout in some way and the web scraper having to be modified to reflect that change
- The retailer adding or removing items from its online offerings
- An issue with the scraping code or server
- The website going offline
- The retailer discontinuing their online shopping service

Figure 2

Number of web-scraped supermarket item price observations per year, broken down by COICOP division



Product churn

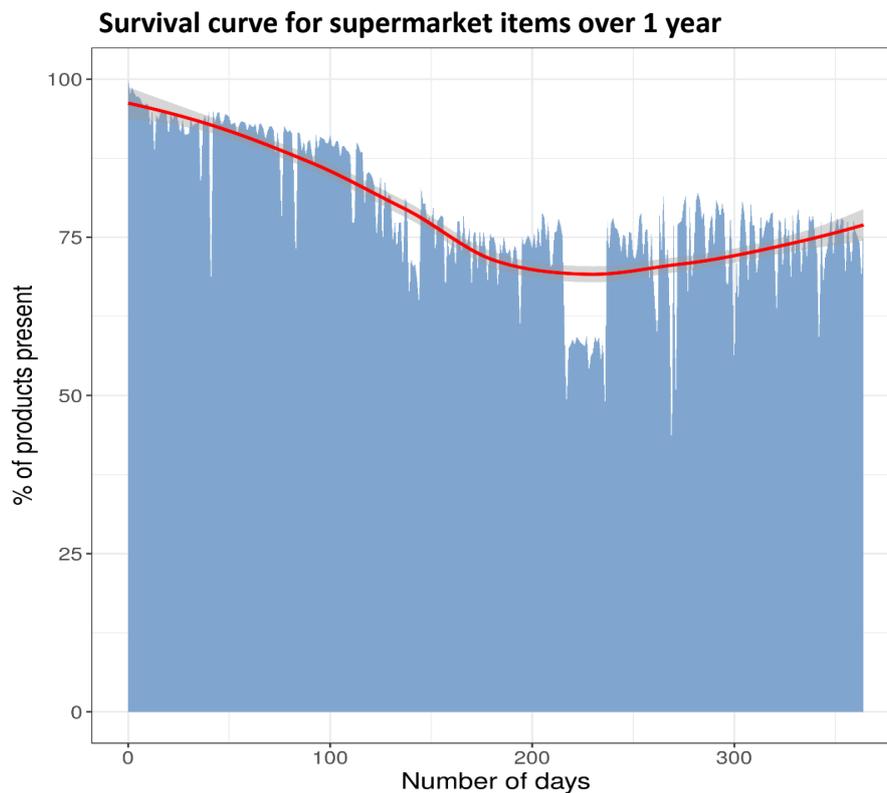
A feature of food prices that lends itself to a TPD chained index is the fact that there tends to be a relatively small amount of churn in food items. Churn is the removal of existing products and introduction of new products. Reasons for product churn may be changing trends or technologies, seasonal changes, clearance (run out items), or relaunching an item.

High churn can lead to bias in a chained TPD. This is because there are fewer items in each window which can contribute to the price index.

Figure 2 shows all the prices scraped from supermarket data over a year. The chosen year had prices from two supermarkets. The figure illustrates the fact that the “Food and non-alcoholic beverage” division does in fact have a low churn in comparison to the other classes in the CPI. Most items in the “Food and non-alcoholic beverage” class are available over 300 days out of a year.

Figure 3 gives a clearer view of the food data, using a survival curve. This data is from the same one-year period as figure 2, with data from both supermarkets. This F curve shows the percentage of supermarket products available on any given day up to year later. Based on this data, 79% of the products available on the first day were also available 365 days later.

Figure 3

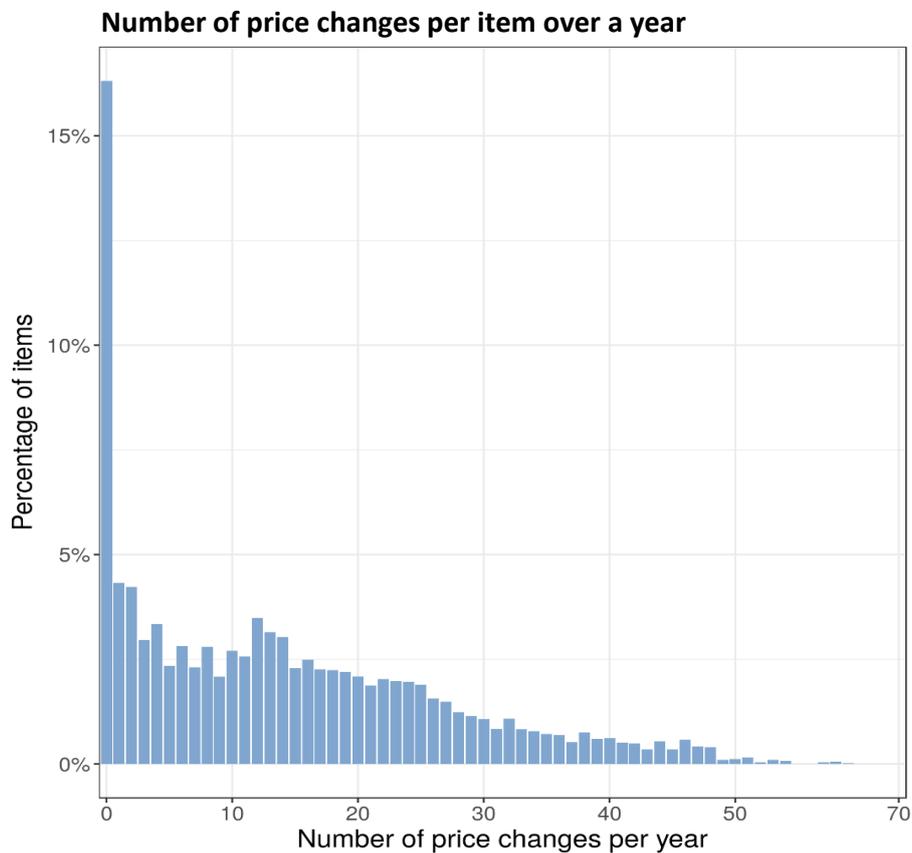


Other features of the web-scraped food prices

Price fluctuations

De Haan and van der Grient (2009) showed that fluctuating prices can lead to a downwards drift in chained indexes, when calculated with expenditure weights. The issue of chain drift will be discussed in [section 5](#). To gain a broad overview of the variability of prices the histogram in figure 4 shows how many products have had a specified number of price changes in a year. Sixteen percent of the products had no change, and 20% had no changes or only one change. From figure 4 we can see another peak in the graph at 12 price changes per year. This is due to many prices having a monthly frequency in price variation. The median number of price changes is 12, and the maximum is 70. Of all items, 99.6% have fewer than 52 price changes in a year – or an average rate of change of greater than one week.

Figure 4



Product classification

As mentioned above in Table 2, the PriceStats dataset is coded at three COICOP levels: division, group, and class. The aggregation of these three levels for food is shown in table 3.

Table 3

Food and non-alcoholic beverages COICOP level 2 and level 3		
Division	Group	Class
Food and non-alcoholic beverages	Food Non-alcoholic beverages	Bread and cereals Meat Fish and seafood Milk, cheese, and eggs Oils and fats Fruits Vegetables Sugar, jam, honey, chocolate, and confectionary Other

To give a rudimentary look into the accuracy of the coding, a random sample of 100 items was chosen from all data in the supermarket set and coded manually. Table 4 shows the accuracy rates we found.

Table 4

Accuracies of COICOP labelling			
Accuracies of labels across entire supermarkets data set			
Level	Division	Group	Class
Accuracy	82%	78%	67%
Accuracies of labels across 'Food and non-alcoholic beverages' division			
Level	Division	Group	Class
Accuracy	100%	100%	88%

While the sample is relatively small the above results indicate the coding accuracies are quite good among the food and non-alcoholic beverages group. Although it is worth noting that the level of aggregation is still quite high. An element of future work will be to build a classification algorithm to classify products down to a lower level. Lower level classification would be useful for drawing out meaningful insights from changes in FPI.

5 Practical considerations for chained TPD indexes

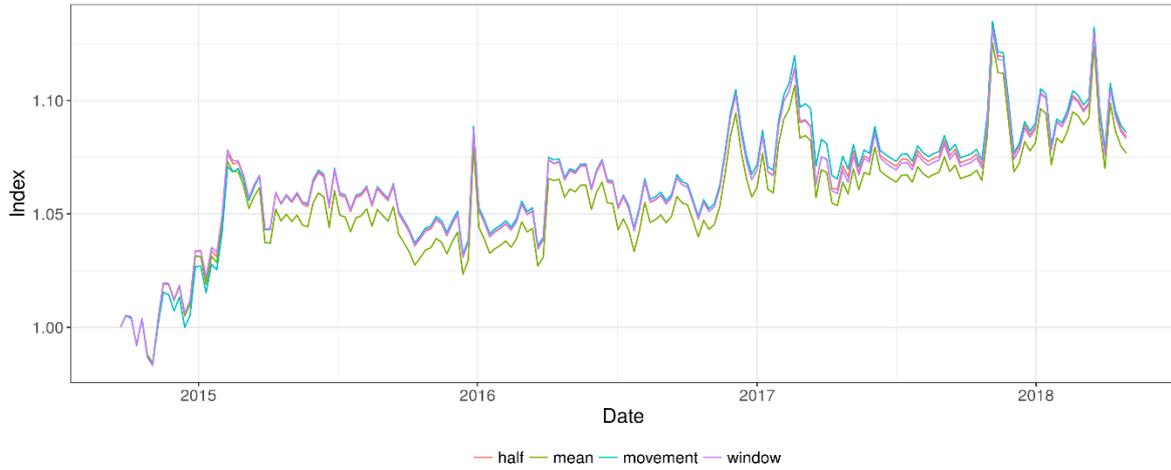
Although the chained TPD does show great promise in calculating an index from web scraped data, some issues must be considered.

Splice position

As discussed in [section 3](#), the literature shows that there are a variety of opinions on the optimal choice of window splice position. As can be seen in figure 5, the overall effect of the splice position is not very significant, so we agree with Diewert and Fox (2017) that the mean splice seems to be a good compromise.

Figure 5

Effect of splice position on the TPD chained index. Data aggregated weekly and by COICOP level 3, window length of 16 weeks.



Window length selection

The selection of window length has a significant effect on the TPD index. Bentley (2018) showed that inflation in rental accommodation can be calculated to be anything from 55% to 127% over a 25-year period, based solely on the selection of window length. The effect of changing window length can be seen in figure 6.

Figure 6

Effect of window length on the TPD chained index. Data aggregated weekly and by COICOP level 3 with mean splice position

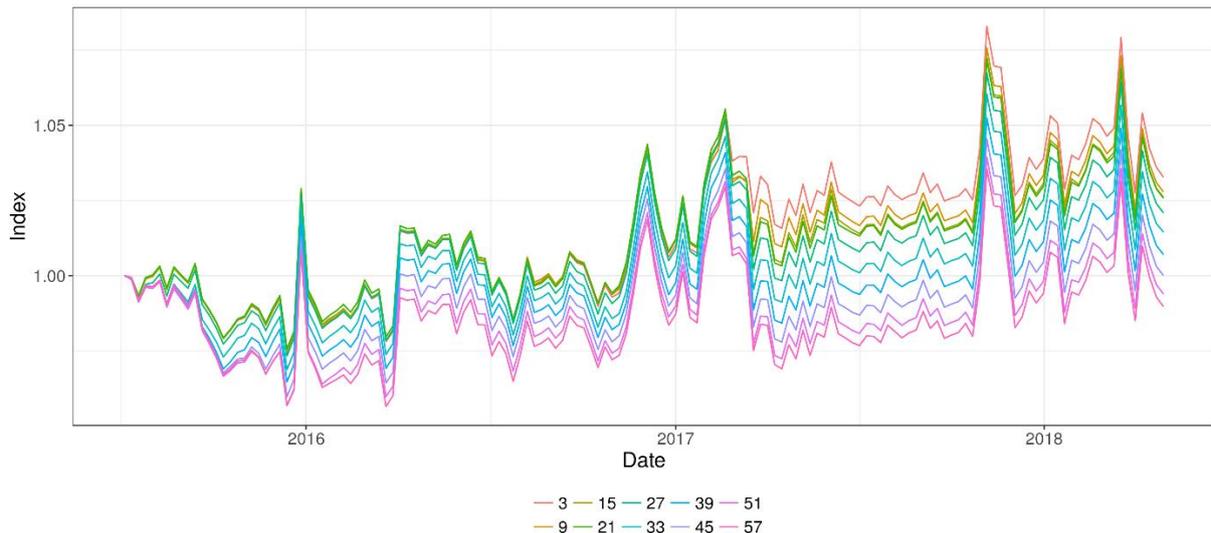


Figure 6 clearly shows that increasing the window length has the effect of reducing the inflation of the index.

In all chained indexes there is a trade-off between the characteristicity and transitivity of the index. Caves, Christensen, & Diewert (1982) described characteristicity as the “degree to which the weights Caves, DW, Laurits, RC, & Diewert, WE are specific to the comparison at hand”. In the case of this data, that is to say that the most recent data is most relevant to the current index, and data from

distant time periods should have less of a bearing on the current index. In other words, decreasing the window length increases the characteristicity.

Transitivity, or circularity, describes the property that “chained indexes will have identical indexes to the corresponding direct indexes”, ILO et al. (2004, p 157). Another way put this is that a highly transitive index calculated quarter on quarter will give very similar results as calculating year on year. Increasing the window length increases transitivity.

As Bentley (2018) points out, these two properties are in conflict, and a compromise must be made that attempts to maximise both properties.

We chose 13 months as a window length as it makes a good compromise between characteristicity and transitivity. This is a common window length choice from the literature (Chessa et al. 2017) as it accounts for items that have a strongly seasonal pattern. This choice of window length may be revised in future as we continue to explore and develop the web scraped data. In addition, there may be a case for considering different window lengths for different sub groups of food.

Price imputation

The approach to dealing with missing data from entire time periods (days, weeks, or months) is another consideration related to the chained TPD index. As discussed in [section 4](#), there is always the chance that there will be discontinuities in web scraped data. Price imputation is only required where there are time periods with no available data at all. There are several options for dealing with this issue. One such method is carrying the last day forward until new data becomes available.

Another method for dealing with missing values is to aggregate the prices over some time period, for example a week or a month for prices collected daily. This will have the benefit of imputing any number of missing time periods, so long as the missing period is shorter than the aggregation time period. An additional consequence of this approach is that time aggregation will decrease the volatility of the index at the expense of losing resolution on the time scale.

At Stats NZ we have not yet had to deal with entire time periods with no data. However, we have aggregated by the month to account for missing time periods should this issue arise.

Chain drift

Chain drift occurs if a chained index “does not return to unity when prices in the current period return to their levels in the base period” ILO et al. (2004, p. 445). Chain drift is an issue commonly observed in chained indexes. The most pronounced cause of chain drift is clearance items. Clearance items are items that are sharply discounted and then removed from stores. They should not be currently used in the CPI as stated by ILO et al. (2004, p. 487).

Clearance items cause chain drift through the following mechanisms:

- As the price drops the index value is reduced.
- When the clearance item drops out, there is no effect to reverse the decrease and the chained index remains suppressed.
- Although each clearance item may only have a slight effect on overall index, the fact that this effect is cumulative causes the chain drift.

Many statistical agencies remove clearance items with the aid of expenditure information from scanner data. Loon and Roels (2018) apply filters such as excluding products with a sharp decrease in sales and in price, and excluding products with a sharp decrease in sales with a constant price to remove clearance items from their calculations.

It is worth noting that Loon and Roels (2018) tested the effect of expenditure weights on their CPI calculations from scanner data. They found that after removing clearance items, the calculation with weights gave very similar results to that without weights, at a high level. The same result was found by van der Grient and de Haan (2011). These are encouraging results for a web scraped FPI which must be calculated without expenditure weights.

Using web-scraped data these options of filtering out clearance items are not available due to the absence of expenditure information. Another approach to combat chain drift is to aggregate the products into groups of items of similar quality. If a group contains an item which does go on clearance, the average price of the group will decrease as the price drops. However, once the clearance item has been removed from the store the group average should return to the pre-clearance price, thus eliminating chain drift.

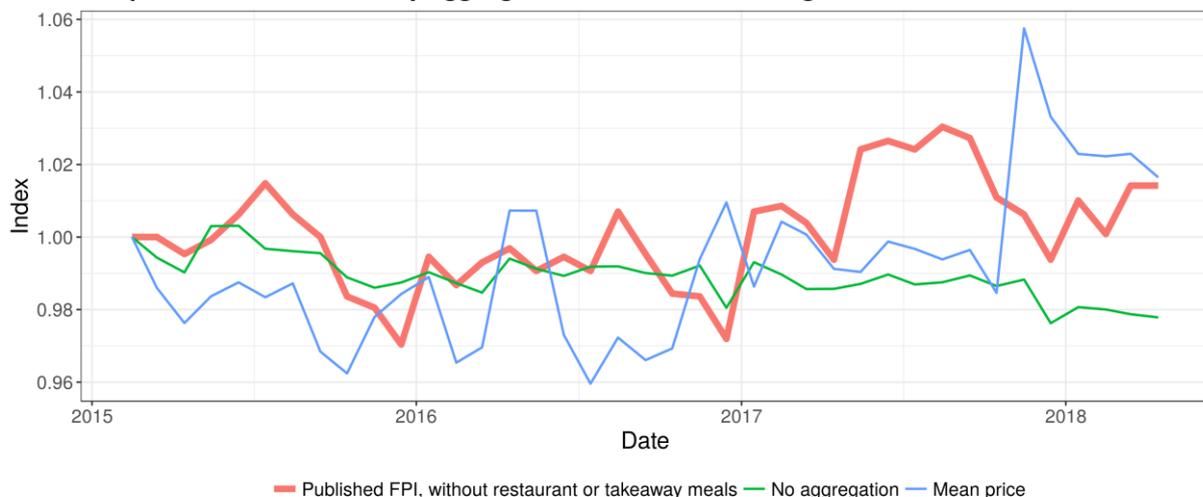
One important point is that by aggregating items, there is an implicit assumption that all items in each group are items of equal quality. By aggregating at too coarse a level, items of differing qualities may be included in the same group. This would violate the constant quality property of the FPI.

6 Results

Figure 7 shows a modified Stats NZ published food price index compared to a chained TPD index calculated using web scraped data with a mean splice and a window length of 13 months. The published FPI includes restaurant and takeaway meals. These have been removed from the FPI in this case to make them comparable to the data contained in the web scraped data. The graph also shows the geometric mean price of all web scraped prices, normalised by the geometric mean of the first time period.

Figure 7

Published FPI less restaurant or takeaway meals compared to FPI calculated from web scraped data with monthly aggregated data, window length of 13 months



No aggregation

Figure 7 shows that a TPD chained model exhibits severe downward drift when calculated without any aggregation or filtering of low expenditure or clearance items. As discussed in [section 5](#), this is mostly due to the presence of clearance items in the dataset.

Mean price

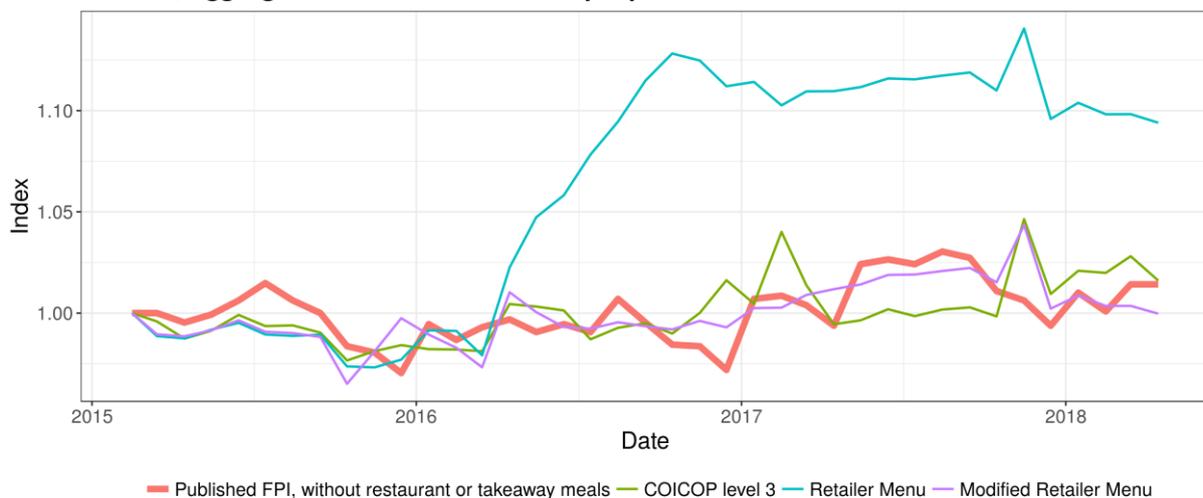
Although the mean price does result in a very similar level of inflation over the total time, the mean price exhibits more variation over the time period. The mean price can be viewed as an index with no adjustment for quality.

Aggregation by COICOP level 3

As discussed in [section 4](#), all items were assigned a level 3 COICOP class by PriceStats. This class, as well as the retailer name, were used to aggregate the prices. The chained TPD was then calculated on this aggregated data with a window length of 13 months and a mean splice. The result, shown in figure 8, is an index that does not exhibit downward drift.

Figure 8

Published FPI without restaurant or takeaway meals compared to FPI calculated from web scraped data with monthly aggregated data, mean splice and window length of 13 months, aggregated based on different properties



However, within this approach there is an implicit assumption that all items from the same retailer and the same COICOP level 3 are of equal quality. If, for example, a retailer were to introduce a new premium range for some items, this approach would result in an inflation of price, and an over-inflated FPI.

Aggregation by retailer menu

In a similar manner to aggregation by COICOP level 3, aggregating by retailer menu also removes downward drift. Note that one retailer changed their retailer menu structure around the first quarter of 2016 resulting in periods with drastic movements. By analysing the retailer menu before and after this change, we devised a concordance to link the two retailer menu structures. We calculated an index from this modified retailer menu. This shows the inherent risks in web-scraped

data, and suggests that statistical agencies would be wise to minimise their dependence on the exact structure or layout of a website.

As with the index calculated by aggregation of COICOP level 3, the retailer menu does not directly allow for quality adjustment of prices. However, as there are more classes in the retailer menu than the COICOP level 3 classification, it has the potential to pick up more subtle changes in products that may reflect quality.

Results summary

In order to compare the differences between the published FPI and each of the indexes presented in figures 7 and 8 the sum of squared differences was calculated for each series. The lower the sum of squared differences between two indexes, the closer the two indexes are to each other.

The sum of squared differences, for an index x , is calculated from the formula below:

$$E_x = \sum_{t=1}^n (P_t^{published} - P_t^x)^2$$

Where:

E_x = sum of squared differences for index x

n = number of available time periods

$P^{published}$ = the Stats NZ published FPI

P^x = the index for which the squared difference is being calculated

Table 5 shows that the modified retailer menu gives the closest fit to the published FPI. Modified retailer menu has enough aggregation to remove the chain drift, but retains enough resolution to account for some quality adjustment. By comparison, the COICOP level 3 aggregation is coarser and so it represents less adjustment for quality.

Table 5

Sum of squared differences of indexes compared to published FPI index		
Index	Sum of Squared difference	Sum of Squared difference, normalised
Published	0	0
Modified retailer menu	0.006	1
COICOP level 3	0.011	1.79
No aggregation	0.015	2.33
Mean price	0.022	2.31
Retailer menu (un-modified)	0.254	39.8

See [Appendix](#) for more detailed graphs when subsetting the data to COICOP level 3.

7 Conclusions and next steps

Conclusion

Work so far has established that it is possible to convert daily web-scraped data into a food price index. This food price index closely matches that of the current approach of sending field officers to supermarkets to collect prices and using expenditure weights (for quantities purchased) from the Household Economic Survey. The web scraped model relies on product aggregation, either by COICOP level 3 or by retailer menu.

Due to the fact that we only have web-scraped data for 3 years, and food has a relatively flat index, it is difficult to say for certain how well this new approach will perform relative to the current method over the long term. In addition, there are other barriers to production as listed below that must be addressed prior to implementing this method in production. However, as a result of this work, we have developed a great deal of knowledge and have identified some other promising avenues to apply the methods discussed in this paper.

Barriers to production

In developing this work, we have observed several barriers to putting this model into production.

Coverage and data reliability

Recently one of the supermarkets in our investigation stopped offering online shopping and moved to an app-based service, which PriceStats are currently unable to scrape. This highlights the high degree of vulnerability in relying on web-scraped data where there is no security of supply. In New Zealand's supermarket sector, we are particularly vulnerable to some stores going offline due to lack of operators. By contrast, there are numerous retailers of electronic goods, so the addition or removal of one store would not constitute a significant change to the overall landscape.

As discussed in [section 4](#), the FPI is of national importance and so the continuity of the index is imperative. If, for any reason, Stats NZ was unable to gather web-scraped data for any given time period it would represent a major failure for the agency. Assurances or contingencies for the absence of web-scraped data are required before we could commit to only using web scraped data, for example:

- Contingency plans include continuing to field collect prices, but only at a limited number of stores.
- Engaging with retailers to see if they can provide some assurances that they plan to continue offering web based shopping, or at least that we will be given advanced warning should they plan to desist.

Quality adjustment

TPD with no aggregation is implicitly quality adjusted. However, by introducing broad aggregation to COICOP level 3 this quality adjustment feature has largely been removed. Bhardwaj et al. (2017) showed UK's Office for National Statistics has developed the CLIP (Clustering Large datasets Into Price Indices) method that clusters like products together based on web-scraped features and aggregates them to calculate an index. A similar approach to this could be developed to categorise like products together, with these groups' categories being fine enough to have an implicit quality associated with them.

Next steps

Stats NZ intends to continue working on “big data” approaches to indexes. Some promising example areas for exploration are:

- Use web-scraped prices to replace field collected prices in the supermarket stores for which this is possible.
- Engage with retailers to discuss more reliable data sources.
- Early investigations into applying the chained TPD to the overseas trade index (OTI) has shown great promise.
- Stats NZ currently purchases some scanner data for other non-food products. We intend to explore using web scraped data and the chained TPD with a view to replace the purchased data.

8 References

- Bentley, A (2018, May). *Rentals for housing: A model-based estimator of inflation from administrative data*. Paper presented at the UNECE Meeting of the Group of Experts on Consumer Price Indices 7-9 May 2018, Geneva.
- Bentley, A. & Krsinich, F (2017, May). *Towards a big data CPI for New Zealand*. Paper presented at the 15th meeting of the Ottawa Group, May 2017, Eltville, Germany.
- Bhardwaj, H. Flower, T. Lee, P. & Mayhew, M (2017). [Research indices using web scraped price data: A summary of the ongoing research into using web scraped price data in the production and development of consumer price statistics](#). Retrieved on 11 June 2018 from www.ons.gov.uk.
- Cavallo, A (2017). Are online and offline prices similar? Evidence from large multi-channel retailers. *American Economic Review*, 107(1), 283–303.
- Caves, DW, Christensen LR, Diewert, WE (1982) The economic theory of index numbers and the measurement of input, output, and productivity. *Econometrica*, 50(6), 1393–1414.
- Chessa, AG, Verburg, J, & Wllenberg, L (2017, May). *A comparison of price index methods for scanner data*. Paper presented at the 15th meeting of the Ottawa Group, May 2017, Eltville, Germany.
- De Haan, J (2015, May). [A Framework for large scale use of scanner data in the Dutch CPI](#) (PDF). Paper presented at the 14th meeting of the Ottawa Group, May 2015, Tokyo. Retrieved from www.stat.go.jp
- De Haan, J, & van der Grient, HA (2009, May). *Eliminating chain drift in price indexes based on scanner data.*, Paper presented at the 11th meeting of the Ottawa Group, May 2009, Neuchâtel, Switzerland.
- Diewert, WE, & Fox, KJ (2017). *Substitution bias in multilateral methods for CPI construction using scanner data* (Vancouver School of Economics Microeconomics.ca working paper – erwin_diewert-2017-3, revised 23 Mar 2017).
- Gibson, A (2017). [First glimpse of New World online shopping](#) (2017, *NZ Herald*, 18 July). Retrieved from www.nzherald.co.nz, on 7 June 2018.
- ILO, IMF, OECD, Eurostat, United Nations, & World Bank (2004). *Consumer price index manual: Theory and practice.*, Geneva: ILO Publications.
- Ivancic, L, Diewert, WE, & Fox, KJ (2009). [Scanner data, time aggregation and the construction of price indexes](#). Retrieved from www.researchgate.net. School of Economics and Centre for Applied Economic Research, University of New South Wales, Sydney.
- Krsinich, F (2011, May). [Measuring the price movements of used cars and residential rents in the New Zealand consumers price index](#) (PDF) Paper presented at the 12th Ottawa Group, May 2011, Wellington, New Zealand. Retrieved from <http://archive.stats.govt.nz>.
- Krsinich, F (2015, July). *Implementation of consumer electronics scanner data in the New Zealand CPI.*, Paper presented at the New Zealand Association of Economists conference, Wellington, New Zealand, July 2015.

Krsinich, F (2016). The FEWS index: Fixed effects with a window splice. *Journal of Official Statistics*, 32, 375–404.

Loon, KV, & Roels, D (2009, May). *Integrating big data into the Belgian CPI*. Paper presented at the UNECE Meeting of the Group of Experts on Consumer Price Indices, May 2018, Geneva.

Van der Grient, HA, & de Haan, J (2011, May). *Scanner data price indexes: The Dutch method versus rolling year GEKS.* Paper presented at the 12th Ottawa Group, May 2011, Wellington, New Zealand.

Waschka A, Milne W, Khoo J, Quirey T, & Zhao, Q (2003). *Comparing living costs in Australian capital cities*. Paper presented at 32nd Australian Conference of Economists, 29 September – 1 October 2003, Canberra, Australia.

9 Appendix

Figure 9
Comparison of published index and the Chained TPD index, broken down by COICOP level 3, mean splice, window length 13 months

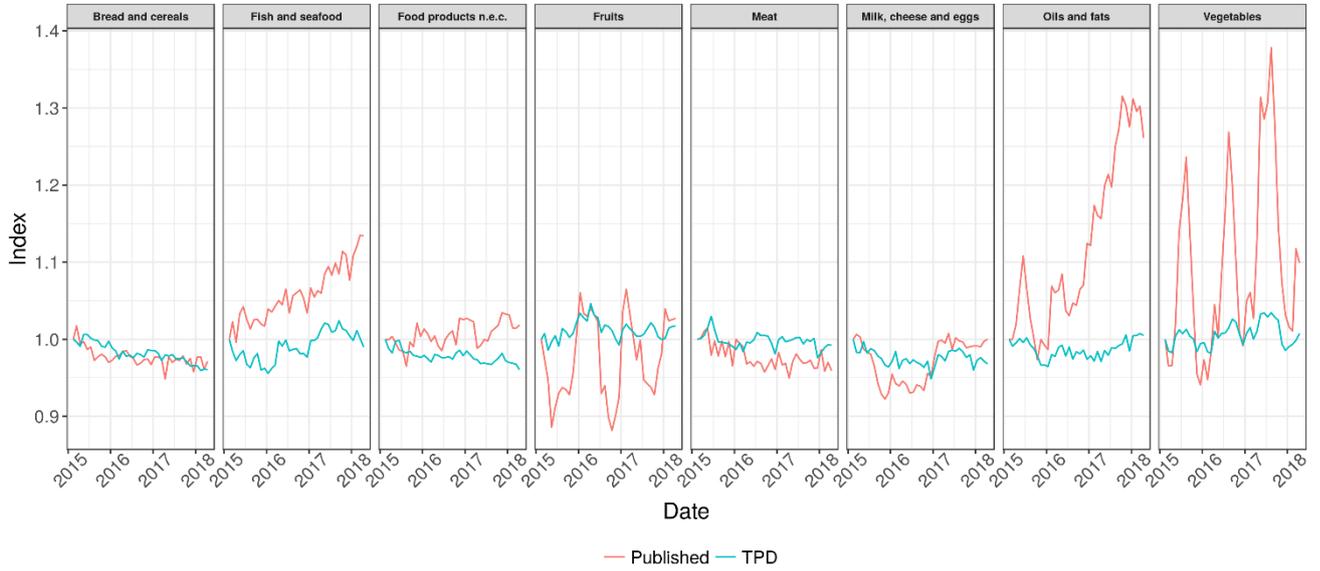


Figure 10
Comparison of published index and the Chained TPD index, broken down by COICOP level 3, aggregated by modified retailer menu, mean splice, window length 13 months

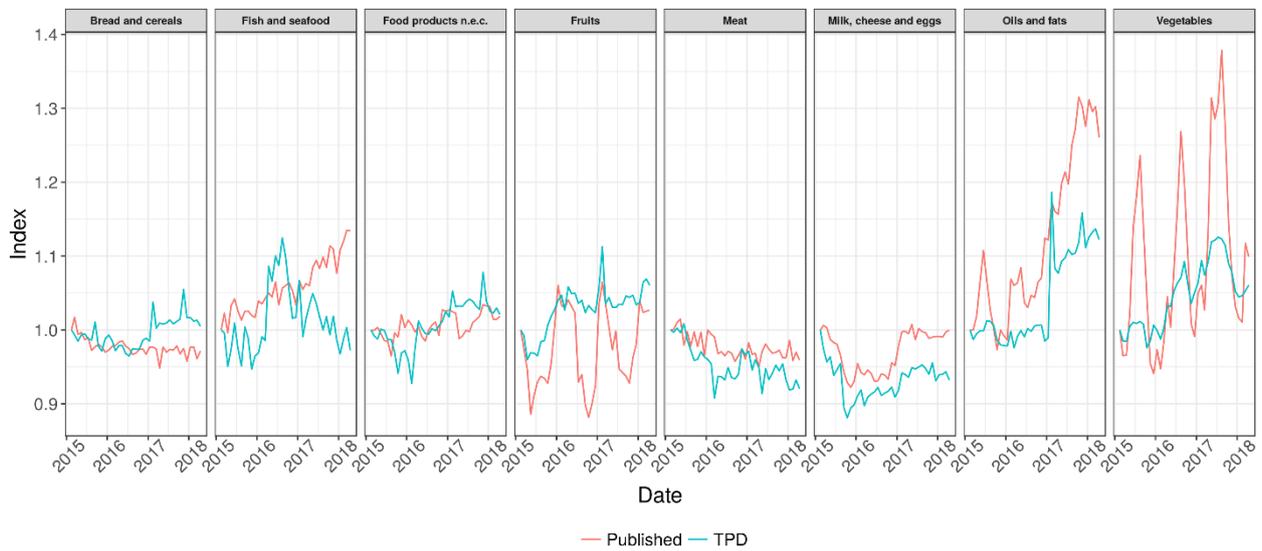


Figure 11

Comparison of published index and mean web-scraped price, broken down by COICOP level 3

