A MAP for the future of price indexes at Stats NZ

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

Introduction

In this paper we give a history of Stats NZ's[1] use of multilateral methods for alternative prices data and explain why we are now developing a generalised research and production in R called the Multilateral Application Pipeline (MAP).

In addition to index estimation, other processes are required in production, and these need to be automated and standardised across different price indexes and data sources to aid transparency, robustness and efficiency. These include:

1. **input diagnostics** to explore and validate source data
2. **output diagnostics** to validate the results of index estimation against those of previous periods
3. **analytical measures** such as decompositions, or 'points effects', to aid insights into the aggregate-level price indexes
4. processes to identify and deal with **changes in the coding of characteristics**, if those characteristics are used for explicit hedonic modelling[2], or if they are required for the creation of unique product identifiers[3]

Because Stats NZ has adopted a range of multilateral methods gradually for a number of data sources over the last 20 years, we now have a range of production processes across SAS, Excel and R, with different levels of automation and robustness. The development of MAP enables us to simultaneously improve our current production processes and pre-build the development and production system for future adoption of new alternative price data sources.

Since Stats NZ has started using these methods, the theory of multilateral price indexes has developed and we are now in a position to develop a system that generalises all the production processes once the source data has been transformed into a consistent format. The appropriate index estimation can then be specified with parameters for each choice of a multilateral index method, a splicing method and an estimation window length, with flexibility to easily change these settings in response to future theoretical findings.

Multilateral price indexes

Traditional index methods do not work well with alternative prices data[4] such as scanner data, administrative data and web-scraped online data for two main reasons:

1. Chained superlative indexes[5] tend to exhibit 'chain drift' when frequent sales result in asymmetric volatility in prices and quantities.
2. Matched-model methods omit the implicit price movements associated with the introduction of new products.

---

[1] Statistics NZ is now called 'Stats NZ'.

[2] For example, in the time dummy hedonic (TDH) or Imputation Törnqvist Rolling Year GEKS (ITRYGEKS) indexes

[3] Such as required for consumer electronics scanner data where model name is masked for those products sold predominantly by one retailer, to protect the confidentiality of that retailer.

[4] Also known as 'non-traditional data' or 'big data' in the context of price measurement, though many argue that these data sources are not strictly 'big data'. A more accurate term might be 'bigger data'.

[5] The seemingly appropriate way to estimate representative price indexes in the context of rapidly changing product universes and full-coverage data.

Over the last 20 years, there has been a significant amount of research and development in this area, resulting in the adoption of multilateral index methods, such as:

- the Time Dummy Hedonic (TDH)
- the Rolling Year GEKS (RYGEKS) (Ivancic, Diewert and Fox, 2011)
- the Time Product Dummy (TPD) (ibid.) or FEWS (Krsinich, 2016)[6]
- the Imputation-Törnqvist RYGEKS (ITRYGEKS) (de Haan and Krsinich, 2014)

Evolution since 2001 at Stats NZ

Stats NZ has used alternative data and multilateral methods in the New Zealand Consumers Price Index (NZ CPI) for **used cars** from 2001; **consumer electronics** from 2014 and **housing rentals** from 2019. In the NZ Overseas Trade Index (OTI), a multilateral method was used for **mobile phones and televisions** from 2013 before being fully adopted **for all price indexes from customs data in the OTI** in 2020.

Krsinich (2014) explains the adoption of multilateral methods at Stats NZ in the wider context of the history of quality adjustment in the New Zealand Consumers Price Index (NZ CPI).

Used cars

Stats NZ first used a multilateral index in production in 2001, when a time-dummy hedonic (TDH) index was adopted as a more efficient and accurate way of estimating price change from a large-scale survey of all used cars sold by a sample of used-car dealers. In 2011 the hedonic formulation was improved and in 2017 administrative data on used cars' characteristics from the New Zealand Transport Authority was incorporated to reduce respondent burden.

Rental prices

In 2009 a time-product dummy (TPD) was used to benchmark the performance of the then matched-model rental index based on a longitudinal survey of landlords. Exploring the properties of this approach then motivated further research by Stats NZ into the potential of using fixed-effects (or time-product dummy) indexes with splicing more generally, for any longitudinal price data with insufficient data on product characteristics to exploit explicit hedonic methods such as the TDH.

In 2019 Stats NZ then redeveloped the rental index in production as a TPD[7] index based on tenancy bond data (Stats NZ, 2019a; Bentley, 2022). The rental price index is currently being redeveloped further by enabling 1. revision to the rental flow index[8] and 2. utilising a 'flash data subset' for provisional estimates to minimise potential revisions bias while maximising the use of data once available. This approach will enable more regional disaggregation, less volatility, more accurate reflection of the timing of price shocks, and a longer historical time-series (Bentley and Krsinich, 2022).

---

[6] The FEWS index explicitly combined window-splicing with a TPD index to address the systematic bias that would result from using a TPD in production for a non-revisable index such as the CPI. Now that splicing (of more than just the latest period) is recognised as an important element in the specification of multilateral methods, the distinction between TPD and FEWS is no longer required and so we will now tend to use the original term 'TPD' to refer to this method.

[7] With a geomean splice.

[8] In contrast to the stock rental index which feeds into the non-revisable CPI.

Overseas trade indexes

In 2013 the TPD was used to estimate price indexes for mobile phones and televisions from import data in the overseas trade index then, in 2020, Stats NZ fully adopted the TPD for estimation of all price indexes from customs data for the NZ OTI (Stansfield, 2019; Stats NZ, 2019b).

Consumer electronics

In 2014 the Imputation Törnqvist Rolling Year GEKS (ITRYGEKS) (de Haan and Krsinich, 2014) index was adopted to estimate price indexes from scanner data for consumer electronics products in the NZ CPI (Stats NZ, 2014).

The future strategy

Bentley and Krsinich (2017) gave an overview of the potential for alternative data in the NZ CPI. Following this, in 2021 an internal review by Stats NZ recommended a strategy for the future of using alternative data in the NZ CPI. The internal report's key recommendation was that Stats NZ should pursue the development of a generalised processing system to consolidate the existing production processes and provide a solid basis for the future incorporation of alternative data sources. The paper by Stansfield and Krsinich (2021) presents some of the conclusions and empirical testing undertaken during that review.

Production processes are non-trivial

At price index conferences and in the literature, most of the focus on the use of alternative data has centred around index methodology. In particular, on the still-evolving concepts, limitations and empirical results relating to multilateral index methods.

However, the estimation of indexes is just one element of what must be dealt with when using alternative data in production. It is also crucial in the production of price indexes to understand 1. what drives aggregate price movements and 2. the impact on the most recent index movement of the splicing procedure used[9]. Issues also arise when dealing with incomplete or inconsistent-across-time source data which, in Stats NZ's experience, is the rule rather than the exception with this data.

Many of the processes required for these insights and mitigations become non-trivial to automate at scale. The iterative development of the MAP system is therefore incorporating an automation of processes which, in the past, have involved relatively time-consuming analytical work, often at-least partially using Excel.

Example: the inconsistent coding of characteristics in consumer electronics scanner data

A good example of an issue arising in production is that the coding of product characteristics can change. A simple example would be the categories (**Y**, **N** and **n.a**) changing to (**Yes**, **No** and **N.A.**)

---

[9] The use of splicing (where the splicing period is greater than just the latest period) trades off the quality of the most recent movement in favour of the longer term index. While this is generally a desirable property, focus is often on the most recent movement (either annual or monthly) so the NSO should understand the impact of the implicit revisioning implied by the splicing.

This type of change happens frequently – in our experience at least a few characteristics have coding changes every quarter[10]. And even if nothing changes, we still need a process that will confirm that.

From the algorithm's point of view, there is no connection between the values **Y** and **Yes** - they are simply different text strings.

There are two reasons we need consistently coded characteristics for estimating price indexes from consumer electronics scanner data. To explain this, first note that we use the Imputation Tornqqvist Rolling Year GEKS (ITRYGEKS) index in production for consumer electronics products. This method is based on a GEKS-Törnqvist (or CCDI (Caves, Christensen and Diewert, 1982; Inklaar and Diewert, 2016) index, with missing values imputed using a time-dummy hedonic (TDH) model based on the pooled bilateral data corresponding to each bilateral Törnqvist index within the estimation window.

The first reason we need consistently coded characteristics[11] is that 'model name' is masked for some of the data[12], so 'pseudo' product identifiers are created based on the combination of values of each characteristic in the data.

The second reason we need consistently coded characteristics is for the hedonic modelling required to impute missing values in the ITRYGEKS index.

With ten different consumer electronics products, with between 10 and 70 characteristics each, identifying and dealing with these potential changes requires at least some automation. In the original SAS and Excel production system, this process was semi-automated but still time-consuming and difficult to understand for new analysts without a deep understanding of multilateral methods.

Figure 1 shows the impact of not identifying and dealing with coding inconsistencies, for an (unnamed) consumer electronics product.

Figure 1. Bias arising from coding changes for a consumer electronics product
Number of matched and unmatched products defined by characteristics' values

---
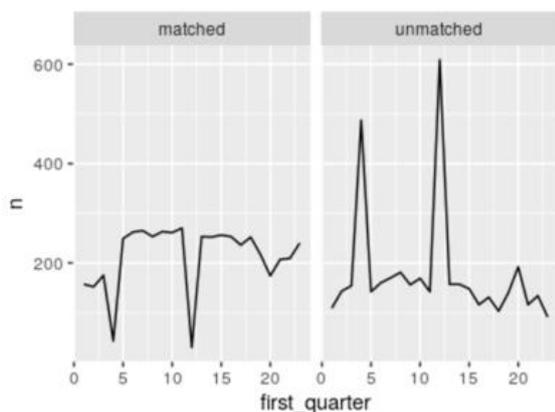
[10] The NZ CPI is a quarterly index.

[11] Note that this would still be an issue with this data if we were using index methods such as the RYGEKS or the TPD which don't require characteristic information for hedonic modelling, as we need longitudinally consistent product identification.

[12] In fact, the data suppliers have now entirely dropped model name from the data for some consumer electronics products.

Quarterly TPD (geomean splice, 5-quarter estimation window) before (full) and after (cleaned) removal of inconsistent characteristics

Index level **Percentage change**
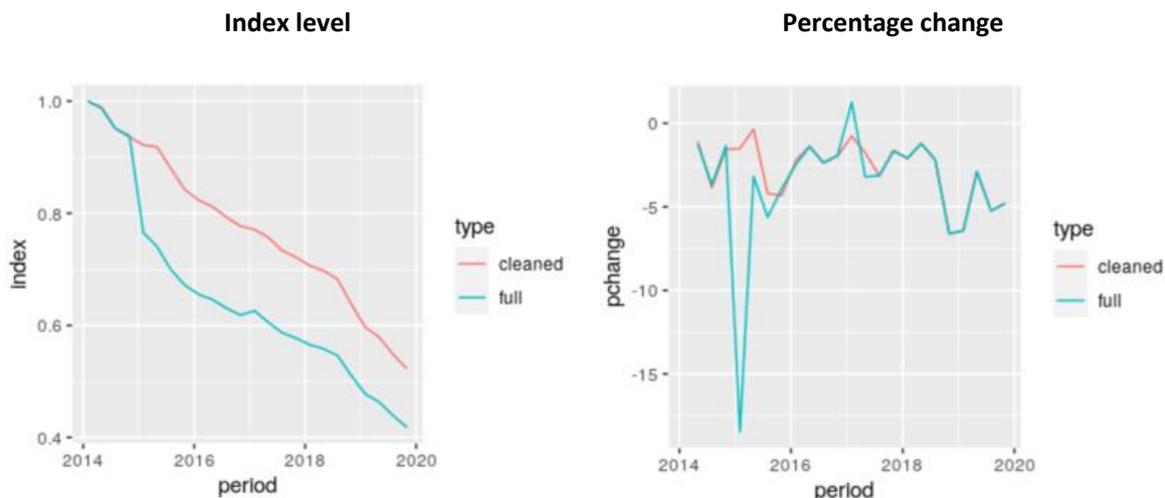
Automating the identification of coding changes

As part of the iterative development of MAP, the identification and treatment of coding changes is being automated to the point where it will no longer require analyst involvement unless the coding change is extensive enough to either require decisions about recoding or to require direct communication and resolution with the data supplier.

The process, which is still being finalised, takes the following approach:

1. Product matching rates are calculated over time within the estimation window[13] of product identifiers created according to the coding of characteristics in the source data.

2. Significant discontinuities (identified using a threshold value, eg 50%) identified in step 1 will then trigger the comparison of number of values of each category of each characteristic across the periods spanning the discontinuity.

3. For each category of each characteristic, if the difference in the number of observations with that category differs across the periods spanning the discontinuity by more than a threshold value, eg 30, then that characteristic is flagged.

4. If more than a number of characteristics are flagged for removal (differing across products such that the remaining number of characteristics is over a certain number) then the process will flag that a major review (possibly involving contacting the data provider) is required.

5. Any flagged characteristics are removed[14] from the derivation of the product id and the hedonic models (if the index method uses hedonic modelling) and the relevant multilateral

---

[13] Multilateral indexes are estimated for window of data which is usually at least a year long. Depending on the index it can be significantly longer (for example the rental index uses an eight-year window).

[14] This might seem extreme, but the number of characteristics in the source data (for consumer electronics) is generally very high with significant collinearity between characteristics. Stepwise regression has found that generally the number of characteristics can be reduced down to under 10 and close to 5 (depending on the product) with little impact on the index (though of course these need to be the right characteristics so in lieu of automating that determination, using all characteristics is the default approach).

index is calculated for comparison with the same index calculated with the original data.

6. If the index on the original data and the data with removed 'problem' characteristics is unchanged (within a very tight tolerance level) for all periods outside those spanning the discontinuity points, then the reduced set of characteristics is maintained for the final estimation.

7. If, on the other hand, removal of the characteristics has an effect on other periods (outside those spanning the discontinuity), a more manual examination of the coding changes will be initiated, with relevant information on the coding changes supplied to the analyst. This will result in recoding of the data which will be versioned within the MAP system such that both the original and recoded data is stored separately and changes documented.

In our testing to date, neither of the flagged processes associated with steps 4 or 7 have arisen but a process to version and document the outcomes of each of these will be incorporated into MAP.

Once this particular production process has been finalised we plan to release it in an R package to accompany the index estimating *multilateral* package described below, alongside other similar production processes built during the development and not already incorporated in the *multilateral* package. These will all integrate with the data structures and formats required by the wider MAP system.

Towards more automation of empirical testing

The ability to automate and scale up both the index estimation and many of the associated production processes is also important when determining the appropriate methods for new data sources. Decisions are required about which underlying multilateral index methods to use (e.g. TDH, GEKS-T, GEKS-IT[15], TPD) and what their appropriate settings should be in terms of splicing method and estimation window length. While some methods will be better than others based on theoretical considerations, we acknowledge that the theory is still evolving. This heightens the importance of empirical testing – across methods and their parameters, and against historical series (where they exist) to help justify those decisions.

For example, to test the feasibility of using online data[16] to produce indexes for supermarket products we used IRI supermarket scanner data with and without expenditure weights. Figure 2 shows that for some products (e.g. blades) weights are very important, while for others (e.g. beer) they are less so.

This has confirmed for Stats NZ that the large-scale use of online data in production of the CPI is undesirable[17], but we are embarking on a collaborative project to test the potential of machine learning models to build indicative inflation indicators from the online data, for use in nowcasting and for forecasting models.

---

[15] The multilateral package refers to ITRYGEKS as GEKS-IT (GEKS Imputation Törnqvist) as a more standardised naming convention. Similarly GEKS-T (or the CCDI) is the Rolling Year GEKS based on a Törnqvist index and GEKS-J is the GEKS based on a Jevons index.

[16] Stats NZ has been purchasing web-scraped online data for a wide range of New Zealand retailers from PriceStats for over six years for research purposes. Krsinich (2015) showed empirical results from testing the use of the unweighted FEWS index (alternatively, the unweighted TPD with a window splice) against the current method used by PriceStats (a chained Jevons) and also simulated the effect of a lack of weights on price indexes for consumer electronics by using scanner data.
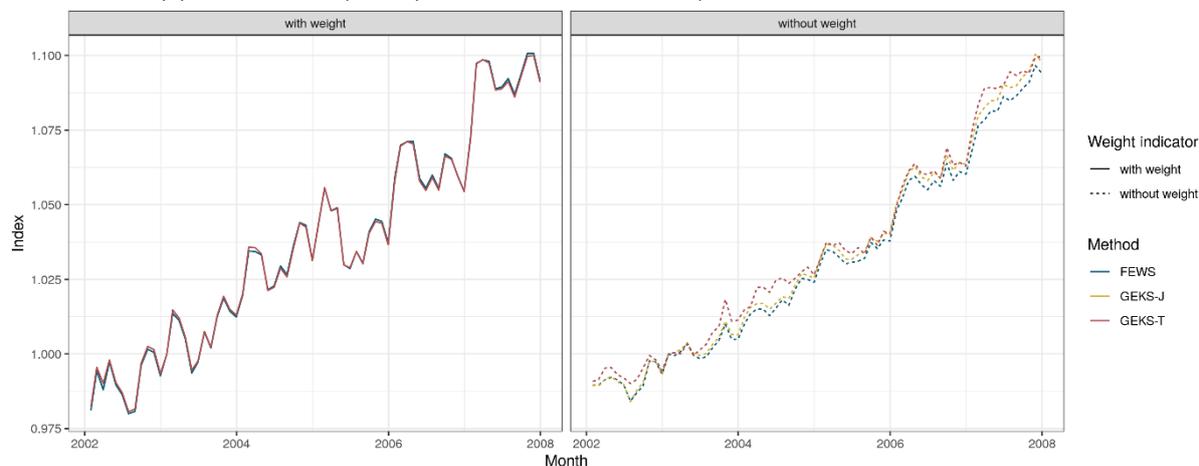
[17] Though we may use it to substitute into the fixed-basket for some products.

Figure 2: Comparing indexes to investigate the importance of expenditure weights

Blades: monthly price indexes (IRI supermarket scanner data)
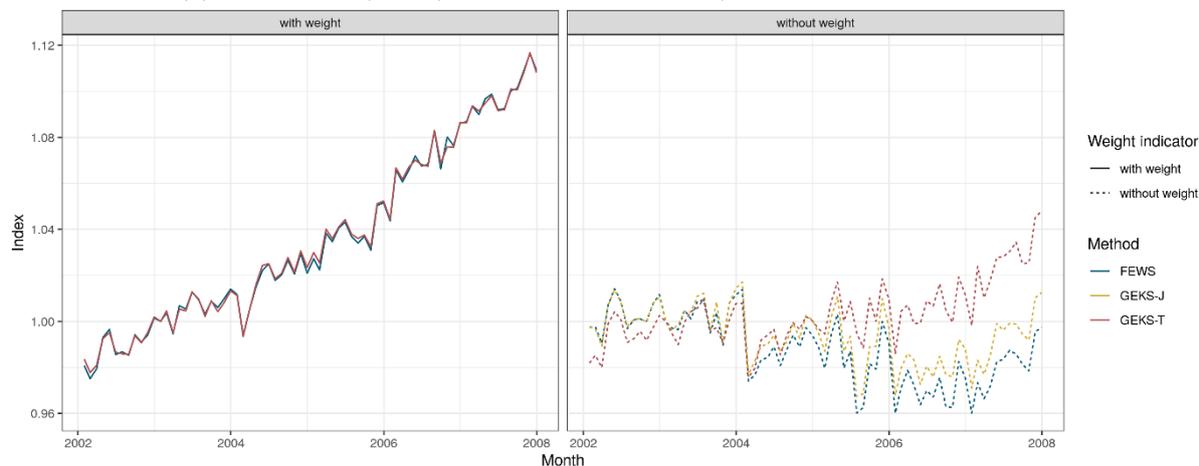
Beer: monthly price indexes (IRI supermarket scanner data)

Stats NZ's internal review included empirical testing for supermarket scanner data using the IRI research scanner data (Bronnenberg, B J, *et al*. 2008). Based on the empirical testing, we are likely to use either TPD (with geomean splice) or the GEKS-T (with geomean splice) once we have the necessary data for all retailers[18] but final decisions will require extensive testing on data for New Zealand, including comparisons to historical indexes.

In addition to the comparison of indexes, as in figure 3, a useful approach we took to comparing the performance of different methods across multiple products was with scatterplots of the correlations of monthly (*x-axis*) and annual (*y-axis*) movements of different methods against some benchmark index. For example, figure 3 shows the correlations against a weighted FEWS[19] index of each of the weighted GEKS-T and the unweighted FEWS, GEKS-T and GEKS-J. For 31 supermarket products we can see that the (weighted) GEKS-T correlates very highly to the (weighted) FEWS. This reassures us

---

[18] Unit values derived from supermarket scanner data have replaced field collection via substitution into the NZ CPI fixed basket since the COVID-19 lockdown, but the full-coverage use of supermarket scanner data with multilateral methods depends on getting full expenditure data from the main NZ retailers, which is in progress.

[19] Actually a FE (or TPD) with a geomean splice – note that Stats NZ has tended to use the term 'FEWS' as a catch-all for TPD indexes with (usually, now) the better-performing geomean splice rather than a window splice.

that theoretical differences between the two approaches are unlikely to be substantive in practice, which makes it a less important decision than it might otherwise be.

Figure 3: Correlations of annual and monthly index movements
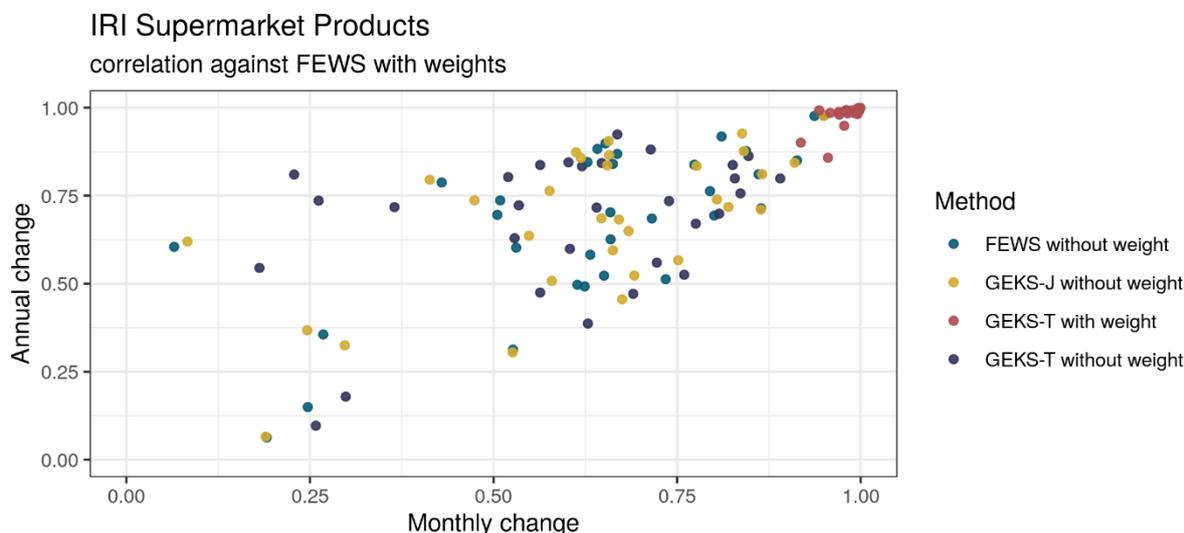
The R packages underlying MAP are open-source

The development of R packages corresponding to the empirical testing described above will be useful in identifying associations between:

- the relative performance of indexes using different methods (and parameters for splicing and window length) and with and without weights, and
- characteristics of the products such as churn rates, price change, and technological change

which seems likely to yield insights that could help guide conceptual development in the ongoing research on the development of multilateral index methods.

By making the R packages underlying MAP open source on CRAN[20], we invite other researchers to consider using them, and potentially further developing them, for research on their own data. Note that the multilateral R package includes a synthetic consumer electronics scanner dataset, which is a useful testbed for development.

The Multilateral Application Platform

As already mentioned, until recently the processing of alternative data sources with multilateral methods at Stats NZ has used bespoke systems across a variety of different languages and operational systems, namely Excel, SAS and R, with varying degrees of manual intervention required by analysts.

---

[20] Stansfield, M (2022).

The earliest implemented processes, such as those for used cars and consumer electronics, are inefficient in various ways by today's standards. For example, the splicing[21] 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. For consumer electronics, the identification and treatment of changed coding for characteristics has been done in excel and is labour-intensive and relatively opaque without documentation of decisions and treatments incorporated into the system itself.

By centralising the process and enforcing 'steady states' in Stats NZ's new Multilateral Application Pipeline (MAP) system, the integration of alternative data sources and multilateral methods can be consolidated and streamlined. Creating a centralised system also brings transparency to these complex processes and a platform upon which team members can learn, with links to documentation and instructions.

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

Multilateral R package for index estimation

Over the past few years we have developed an R package for estimating all the multilateral packages in production at Stats NZ, the *multilateral* package, which is now available at CRAN[22]. Some of the underlying functions are an implementation of the *IndexNumR* package[23] by Graham White. We have also added multilateral methods that use hedonic regression modelling, such as the time dummy hedonic (TDH) and the Imputation Törnqvist Rolling Year GEKS (ITRYGEKS[24]).

Stats NZ built our own package internally to ensure full transparency, particularly for our validation against existing SAS-based implementations, and with consideration of speed and the flexibility to change between methods and parameters easily. For speed of processing, the package allows parallel processing and optimized functions like sparse matrices and memory efficient operations. The extra hedonic regression functionality, in particular, is computationally intensive and requires this optimization.

Figure 4 shows the relative processing times within the Stats NZ environment using parallel processing (with four CPU cores) compared to standard runs (one CPU core) on two years of data of approximately 50 million observations. Both the GEKS-T and TPD methods use geomean splicing and an estimation window length of 13 months.

---

[21] The CPI is non-revisable. Multilateral methods, however, re-estimate a back series with each successive period. This means that new results must be 'spliced' onto the published index such that they preserve the integrity of the published series (by incorporating a 'revision' factor). See de Haan (2019) for a discussion of different splicing approaches.

[22] Comprehensive R Archive Network hiips://cran.r -project.org/web/packages/multilateral/index.html

[23] hiips://mirrors.pku.edu.cn/cran/web/packages/IndexNumR/index.html

[24] Referred to as GEKS-IT in the *multilateral* R package.

Figure 4: The effect of parallel processing on run-time (in minutes) of the GEKS and TPD indexes
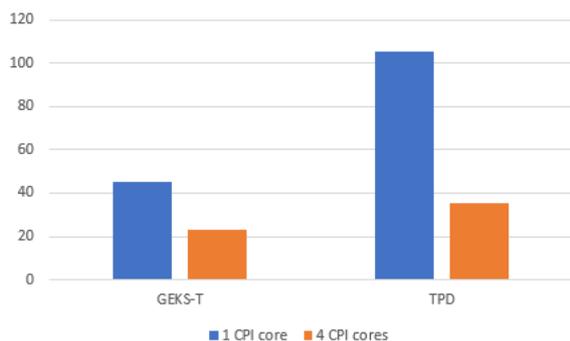
GEKS-T *45 min (1 core), 23 min (4 cores)* **TPD** *105 min (1 core), 36 min (4 cores)*

The *multilateral* R package is the index-estimating R package that sits within the wider Multilateral Application Pipeline (MAP) system.

Overview of the wider MAP system

The current iteration of the MAP system consists of six core processes:

01 Environment setup: where the user can observe key variables and make intentional changes, for example - the period of calculation.

02 Folder creation: this step automatically builds the environment and nested structure for all outputs relating to the production period.

03 Outlier detection: a generalised solution to outlier detection adapted from the BANFF system (Kozak, 2005).

04 Calculate: the index estimation and associated diagnostic and decomposition measures.

05 Format for production: this step takes the standardised output of the calculation and formats it ready for the production system.

06 Reporting: this step takes all the diagnostic files generated throughout the process and generates a markdown (pdf) file for analysts to interpret the input and output diagnostics.

These core processes are supported by a range of generalised helper functions and specific topic[25] supporting functions. Figure 5 gives a visual representation of the MAP system. The core processes are initialized with parameters and options specific to each topic, specifying choices such as the index methods, splicing methods, estimation window length, relevant diagnostics, and other production processes (such as consistency checks for characteristics as described above).

---

[25] A 'topic' refers to a data source or product group. For example, consumer electronics and used cars are two separate topics.

Figure 5: overview of the MAP system
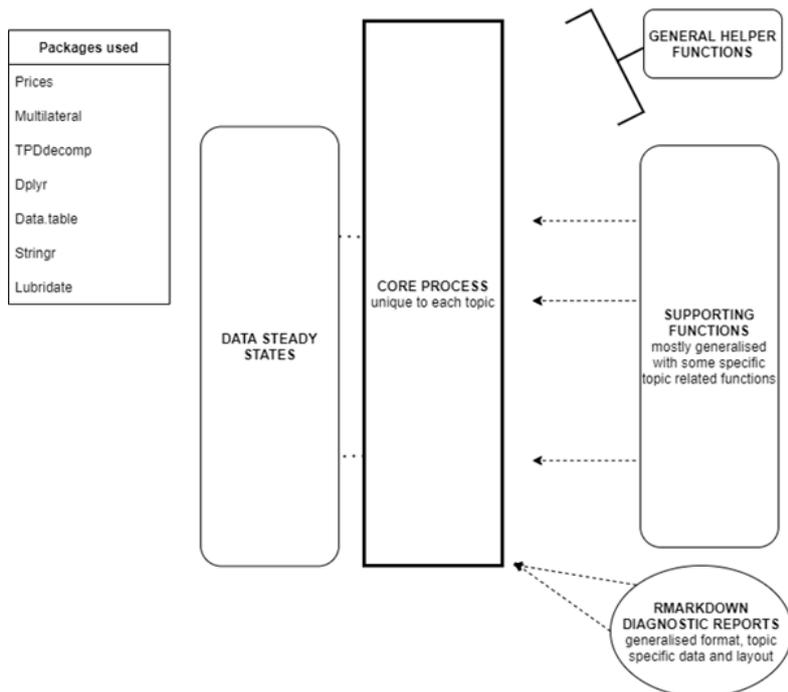
The iterative development of MAP

An *Agile* approach has been taken in the development of MAP, migrating one existing system at a time and identifying new issues as they arise and updating the system in response. Parallel runs of each existing system and MAP are undertaken to ensure exact replication is achieved. Because the two systems are completely independent from the source data forward, a rigorous check of internal processes is ensured.

Example: migrating the consumer electronics scanner data system

The original production system for consumer electronics scanner data was implemented in 2014 in SAS, with diagnostic and analysis processes largely executed in Excel. The system required manual intervention from analysts to produce and respond to diagnostic processes. At the time it was introduced, the system was well understood but with staff turnover the process has slowly turned into a 'black box' with gradual loss of understanding of the purpose underlying key steps. This has made the system quite fragile and overly dependent on a few senior technical staff to deal with ad-hoc issues.

The system usually took about four days for an analyst to run[26], as issues generally arise and require bespoke problem-solving. If no issues at all arise the fastest possible run (which incorporates some quite laborious semi-automated work in excel) takes approximately 4 hours.

With its migration to MAP this system can now run in one hour end-to-end, with all reports automatically produced. Generally no manual intervention is now required in response to diagnostic processes. Version control also makes the process much easier to maintain.

---

[26] While in today's terms this is a long time, the field-based process that scanner data replaced took weeks, as it required time-consuming in-office quality adjustments by analysts of field-collected prices.

Future plans to migrate into MAP

To date, used cars and consumer electronics have been migrated to MAP. The next systems to migrate will be the rental price index (which use tenancy bond data) and the overseas trade indexes (which use customs data). Although these already have their own relatively robust R-based systems, they will be rebuilt in the generalized MAP system to enable full consolidation and streamlining.

Likely future data sources to be developed in the MAP system are:

- **Supermarket scanner data**, once we gain full expenditure data from retailers
- A **prototype official house price index** able to be disaggregated into land and structure indexes, using local councils' valuation and sales data (see Krsinich, 2019)

We are also now exploring the potential to use MAP inside the NZ GS1[27] environment to produce indexes securely with release to Stats NZ of the aggregate-level indexes. This is looking very promising.

Conclusion

In addition to the methodological challenges of using alternative data for price index estimation, there are non-trivial issues associated with production at scale. Our development of the R-based Multilateral Application Pipeline (MAP) helps to automate, consolidate and generalise these production processes.

The development of MAP has been iterative, starting with the migration of existing production systems for used cars and consumer electronics products, from SAS and Excel. Later this year Stats NZ will migrate the more recently-developed R-based systems for the Overseas Trade Index (based on customs data) and the Rental Price Index (based on tenancy bond administrative data).

We plan to develop supermarket scanner data and a prototype house price index using the MAP system, and we are currently exploring the use of MAP inside NZ GS1's secure environment to enable the safe use of confidential price and quantity data linked to barcode information.

For Stats NZ, there are multiple benefits of the MAP system:

- A reduction in manual, error-prone processes – everything that can be automated will be automated.
- More transparency, with the underlying code open for review and reuse by others.
- Diagnostics, monitoring and analysis are incorporated alongside index estimation.
- Index estimation is done with our in-house developed *multilateral* R package, which enables the full range of multilateral methods already in production at Stats NZ, and performs well at scale through optimised functionality and parallel processing.
- Consistent interfaces and processes across product types, data sources and methods.
- The potential for incorporation of links to training and documentation in the user interface.

With MAP's R packages open source and available from CRAN, we hope that other agencies and researchers will also make use of them in their research and development.

---

[27] GS1 hold price and quantity information corresponding to their barcode information from a market research company, meaning that sufficient information for (non-hedonic) multilateral index (such as TPD or GEKS) methods is available within their secure environment, though not able to be released at that level of disaggregation.

References

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

Bentley, A (2022) *Rentals for Housing: A Property Fixed-Effects Estimator of Inflation from Administrative Data* Journal of Official Statistics, 38(1)

Bentley, A and Krsinich, F (2022) *Timely Rental Price Indices for thin markets: Revisiting a chained property fixed-effects estimator* Paper presented at the 2022 Ottawa Group conference, Rome, Italy

Bronnenberg, B J, M Kruger and C F Mela (2008) *The IRI Marketing Data Set* Marketing Science 27: 745-748

Caves, D W, L R Christensen and W E Diewert (1982) *The Economic Theory of Index Numbers and the Measurement of Input, Output, and Productivity* Econometrica 50, 1393-1414

de Haan, J and Krsinich, F (2014) *Scanner data and the treatment of quality change in nonrevisable price indexes* Journal of Business and Economic Statistics, 32(3)

de Haan, J (2019) *Rolling Year Time Dummy Indexes and the Choice of Splicing Method* Paper presented at the 2019 Ottawa Group conference, Rio de Janeiro, Brazil

Inklaar, R and W E Diewert (2016) *Measuring Industry Productivity and Cross Country Convergence* Journal of Econometrics 192, 426-433

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

Kozak, R (2005) *The BANFF system for automated editing and imputation* Proceedings of the Survey Methods Section, 2005 SSC Annual Meeting

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 (2015) *Price indexes from online data using the fixed-effects window-splice (FEWS) index* Paper presented at the 2015 Ottawa Group conference, Tokyo, Japan

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

Krsinich, F (2019) *Land prices: UNCOVERED! Extricating land price indexes from improved property price indexes for New Zealand* Paper presented at the 2019 New Zealand Association of Economists conference, Wellington, New Zealand

Stansfield, M (2019) *Import and export price indexes using fixed-effects window-splicing* Paper presented at the 2019 New Zealand Association of Economists conference, Wellington, New Zealand

Stansfield, M and F Krsinich (2021) *Bigger, better, faster: further progress in using non-traditional data to measure price inflation* Paper presented at the 2021 New Zealand Association of Economists conference, Wellington, New Zealand

Stansfield, M (2022) *Multilateral R package* available on the Comprehensive R Archive Network (CRAN)

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*