Scanner Data Price Indexes: The "Dutch Method" versus Rolling Year GEKS

Heymerik A. van der Grient[a] and Jan de Haan[b]

March 2011

Abstract: In 2010, Statistics Netherlands expanded the use of scanner data in the CPI to six supermarket chains. At the same time a new method was introduced to construct the scanner data indexes. Monthly chained Jevons price indexes are now computed at the lowest aggregation level. This approach was chosen to maximize the number of matches in the data through chaining while circumventing the problem of chain drift. However, the lack of weighting at the item level is an obvious weakness. The solution proposed by Ivancic, Diewert and Fox (2009) is to apply (a rolling version of) the GEKS method, known from spatial price index measurement, to intertemporal price measurement. To make an informed decision about possible implementation, Statistics Netherlands has a shadow system running which computes rolling year GEKS indexes for each COICOP category and each supermarket chain and aggregates the indexes up. This paper briefly outlines both methods and presents monthly index numbers for 2009 and 2010. At the all items, all chains level the two methods yield similar results. At lower aggregation levels there appear to be some marked differences.

Key words: chain index, consumer price index, GEKS method, scanner data.

---

[a] Division of Macroeconomic Statistics, Statistics Netherlands; ha.vandergrient@cbs.nl.

[b] Corresponding author; Division of Macroeconomic Statistics, Statistics Netherlands, P.O. Box 24500, 2490 HA the Hague, The Netherlands; j.dehaan@cbs.nl.

Paper prepared for the twelfth meeting of the Ottawa Group, 4-6 May 2011, Wellington, New Zealand. The views expressed in this paper are those of the authors and do not necessarily reflect the views of Statistics Netherlands. The authors are grateful to Wim Kiewiet for developing the system that computes RYGEKS index numbers.

1. Introduction

The international CPI Manual (ILO *et al*., 2004) notes that "Scanner data constitute a rapidly expanding source of data with considerable potential for CPI purposes" (p. 54); "Scanner data obtained from electronic points of sale include quantities sold and the corresponding value aggregates on a very detailed level" (p. 92); "Scanner data are up to date and comprehensive" (p. 478). In spite of the potential advantages, as far as we know only three countries, i.e., Switzerland, Norway and the Netherlands, are actually using scanner data in the compilation of the CPI; see also Ivancic (2010).[1]

Scanner data can be exploited in different ways. The simplest way would be to use scanner data as an alternative source for price collection, which replaces collection in the stores, without changing the traditional principles of computing the price indexes. More specifically, the samples of items observed and the expenditure weights would be left unchanged. This method is currently applied by the Swiss Federal Statistical Office (Becker-Vermeulen, 2006). In the Netherlands and Norway a more rigorous approach has been followed, though the computation methods differ. The sub-index for food and non-alcoholic beverages in the Norwegian CPI, introduced in August 2005, is based on chained Törnqvist indexes at the lowest level of aggregation (Rodriguez and Haraldsen, 2006). The Dutch scanner data indexes for supermarkets were introduced into the CPI already by June 2002. They relied on annually chained Laspeyres indexes at the lowest level (Schut, 2002). In Norway and the Netherlands, both prices and weights for a large sample of items from a product category are derived from the scanner data.

In January 2010 Statistics Netherlands expanded the use of scanner data (van der Grient and de Haan, 2010). Seven more supermarket chains have been found willing to co-operate and regularly supply scanner data but two of them were not yet implemented in January 2010.[2] The six chains for which scanner data were included in January 2010 have an aggregate market share of some 50% and a weight of slightly over 5% in the CPI.

---

[1] A number of statistical agencies are investigating the use of scanner data at the moment, in particular the Australian Bureau of Statistics, the U.S. Bureau of Labor Statistics and Statistics New Zealand. In the past, several statistical agencies around the world have been looking into the issue. There is also a large academic literature on scanner data, both on price index construction, including quality adjustment, and on broader economic measurement issues.

[2] Scanner data from one of those two chains was introduced in January 2011. It is anticipated that the data from the other chain will be introduced in the near future.

Apart from increasing the accuracy of the CPI (and the European Harmonized Index of Consumer Prices, HICP), there are two more reasons behind the wider use of scanner data in the Netherlands: to raise cost efficiency and to lower response burden. Visiting supermarkets to collect shelf prices is a major cost component of producing a CPI in the traditional way. A total reduction of around 15,000 price quotes each month has been attained for the six supermarket chains.[3] The use of scanner data is beneficial to the data providers as well. They will no longer be bothered by price collectors who walk around in the stores collecting prices or asking staff for help. Lowering response burden is a key issue for Statistics Netherlands.

The expansion of the use of scanner data in January 2010 was accompanied by a new index construction method at the lowest aggregation level. While expenditure data is available for individual items, this information is only indirectly used, i.e., to select a sample of items for a particular product category through cut-off sampling. No explicit weighting is applied in the index computation; monthly chained unweighted geometric, or Jevons, price index numbers are calculated. This method was chosen to maximize the number of matches in the data through chaining – which was deemed necessary given the high attrition rate of items in scanner data – without running the risk of introducing chain drift. As was demonstrated in a number of empirical studies (see e.g. Feenstra and Shapiro, 2003; Ivancic, Diewert and Fox, 2009; and de Haan and van der Grient, 2009), high-frequency chaining of weighted price indexes, including superlative indexes such as the Fisher and Törnqvist, can lead to huge chain drift.

Yet the lack of weighting at the item level is a weakness of the Dutch method and has been criticized on various occasions. Ivancic, Diewert and Fox (2009) proposed an appealing solution to the problem of chain drift in superlative indexes. They adapted the (G)EKS approach from international price comparisons to price comparisons across time. To address the problem of revisions, they suggested a rolling year (RY) version. During 2010, the RYGEKS method was extensively analysed at Statistics Netherlands. In order to make an informed decision about possible implementation, a shadow system was developed that computes RYGEKS indexes on a monthly basis. This enabled us to make a full-scale comparison with the currently computed index numbers. The present paper presents a selection of the results.

[3] It should be noted that the initial number of price quotes for supermarkets was relatively large given their weight in the CPI.

The paper is structured as follows. Section 2 gives an overview of the Dutch method for treating supermarket scanner data. Section 3 briefly outlines the RYGEKS method. Section 4 describes the scanner data used and a number of activities, like data cleaning, which are carried out before actually computing the indexes. Section 5 then compares the index numbers based on the two methods. Section 6 discusses issues that arise in practice when working with scanner data. Some of these issues apply in general, regardless of the computation method, whereas others are method dependent. Section 7 concludes by describing the advantages (and a few potential practical drawbacks) of the RYGEKS approach over the currently used method.

2. The "Dutch Method"

2.1 An Overview

Scanner data reveals that purchases in supermarkets are highly dynamic. For example, for some product categories the attrition rate is huge: many "old" items disappear and many new ones appear. A fixed-basket index would ignore these dynamics. This effect can be mitigated by keeping the sample up to date when items leave but, depending on the sample size, this might mean a lot of work for the price statisticians. This was one of the reasons for Statistics Netherlands to seek for an alternative method. It also happens that items temporarily disappear from the data sets, either because they are out of stock or because purchases are zero.

Scanner data also reveal the distribution of the item expenditures in each month. There are two important aspects. First, Dutch consumers seem to react instantaneously to promotional sales. De Haan (2008) presented an example on detergents showing that the quantities purchased increase enormously when the price is reduced.[4] The fact that prices of supermarket products are not sticky at all has been observed and analyzed in many academic studies. This has the implication that, at low aggregation levels, scanner data price indexes could be rather volatile. Second, even without sales, the distribution of expenditures within a product category is usually highly skewed. That is, a relatively

[4] Triplett (2003) argued that this could be partly due to consumers who mainly purchase the goods when they are on sale. If this is the case, then the population of consumers that buy at a particular store or at a particular supermarket chain will continuously change over time.

small number of items account for the majority of expenditures. To put it differently: in many cases, 50-60% of the items account for less than 10% of the total expenditures. Thus, the use of an unweighted index number formula would overstate the importance of the many low-expenditure items.

Given the high number of exits and entries observed in scanner data, the choice made by Statistics Netherlands for a chain index at the lowest aggregation level seems quite natural. As mentioned in the introduction, chaining of weighted indexes can lead to substantial chain drift, so the unweighted Jevons index formula was chosen. To take the skewed expenditure distribution into account, a crude type of *implicit* weighting was introduced through cut-off sampling: the most important items in terms of expenditures are included in the sample with certainty whereas less important items are excluded.

More specifically, an item is used in the computation of the price index between two adjacent months if its average expenditure share in the current and preceding month with respect to the set of matched items – those items that are purchased in both months – is above a certain threshold value. The threshold was chosen such that roughly 50% of the items is selected, representing 80-85% of aggregate expenditure. This was done at the elementary aggregation level, i.e., for each product category at the most detailed level within each supermarket chain.[5] Our price concept is the unit value (expenditures divided by the quantities sold) for an item within a supermarket chain. So we aggregate across stores belonging to the same chain but not across different chains.[6]

A drawback of an adjacent-period matched-items approach is that temporarily unobserved items are excluded from the computation. This means that the price changes of those items occurring between the last month they were in the sample and the month they re-enter the sample would be ignored. The “missing prices” are therefore imputed by multiplying the last observed price by the (Jevons) price index of the matched items within the same elementary aggregate, as usual. In a way a panel element is forced onto the matched-items approach so that price changes occurring during a period of absence are included in the index.

[5] We use the term elementary level for the aggregation level below which expenditure information is not used in the Dutch approach, which coincides with the lowest level at which price indexes are constructed (though not published).

[6] Ivancic and Fox (2010), using Australian scanner data, found that items can be deemed homogeneous across stores belonging to the same chain. This means that aggregating across those stores makes sense to compute unit values. We assume that their results also apply to the Netherlands.

Like in any matched-items method, quality changes are not explicitly taken into account. Implicit quality-adjustment methods have been most prominent in the Dutch CPI in the past, and in this respect the new method is comparable to former practices. The computer system does allow for making explicit quality adjustments, in particular adjustments for changes in package and pack size, but we expect this feature to be used infrequently.[7]

The aggregation of elementary price indexes computed from the scanner data is similar to the aggregation of price indexes computed from other sources. This is based on annual chaining of fixed-weight ("Laspeyres-type") indexes, where the weights refer to the previous year and where short-term index series are chained in December. The elementary aggregation level can be seen as a sub-division of the 5-digit COICOP level, referred to as "6-digit COICOP". It is constructed using product classifications provided by the supermarket chains, which makes it chain specific. The advantage of using these classifications is that Statistics Netherlands does not need to classify items into product categories, an activity which previously proved to be very labor intensive.[8]

Because the elementary ("6-digit COICOP") level is chain-specific, the first step of the aggregation procedure is to aggregate the elementary price indexes for each chain to the official 5-digit COICOP level and to higher levels. At each of these COICOP levels the price index numbers are subsequently aggregated across the different chains. The weights are derived from the scanner data. This detailed weighting information was unavailable prior to the use of scanner data. Not only do the aggregate market shares of the supermarket chains differ substantially, which we already knew, they also appear to differ significantly across product categories.

The computation and aggregation of scanner data indexes is a separate module in the automated CPI process. The final aggregation step is to combine the scanner data indexes at all COICOP levels with the corresponding indexes estimated from manual price collection in other store types. To this end a so-called matrix of expenditures was developed, which is updated every year.[9]

[7] Fox and Melser (2011) address the issue of adjustments for changes in pack and package size.

[8] A limited number of items may not be classified properly and need to be re-classified.

[9] Although the availability of expenditure data that can be used as weights is a big advantage, at the same time it may reveal inconsistencies with data obtained from other sources used to construct CPI weights such as the national accounts, the household expenditure survey, etc. Adjustments are then necessary to construct a consistent overall weighting scheme.

2.2 The Actual Calculation

In this section we describe the computation of the scanner data indexes for a particular supermarket chain in greater detail. Data cleaning procedures, which precede the actual computation, will be addressed in section 4. The following notation is used. The price (unit value) and expenditure share of item *i* for elementary aggregate *a* in month *m* of year *y* are denoted by $p_{i,a}^{y,m}$ and $s_{i,a}^{y,m}$;[10] $N_a^{(y,m-1),(y,m)}$ denotes the number of matched items between months *m* and *m*-1 of year *y* (where *y*,*m*-1 is equal to *y*-1,12 for *m*= 1).

To introduce a crude type of weighting, every item *i* is given a probability $w_{i,a}^{y,m}$ to be included in the sample used to compute the price change between month *m*-1 to month *m*. These inclusion probabilities or implicit weights are given by

$$w_{i,a}^{y,m} = 1 \qquad \text{if } \frac{s_{i,a}^{y,m-1} + s_{i,a}^{y,m}}{2} > \frac{1}{N_a^{(y,m-1),(y,m)}\chi};$$

$$w_{i,a}^{y,m} = 0 \qquad \text{otherwise.}$$

This means that if the item's average expenditure share in months *m*-1 and *m* exceeds the threshold $1/(N_a^{(y,m-1),(y,m)}\chi)$, the item will be included in the sample. Notice that the sum of all implicit weights determines the sample size of elementary aggregate *a*, that is, $\sum_{i=1}^{N_a^{(y,m-1),(y,m)}} w_{i,a}^{y,m} = n_a^{(y,m-1),(y,m)}$.

The parameter χ must be set in advance. It can be given any positive value, but in practice there is a lower boundary because if the value is too low, the sample will be empty. We decided to use $\chi = 1.25$ for all elementary aggregates, i.e., for all product categories at the pseudo 6-digit COICOP level and all supermarket chains. A simulation study did not indicate an urgent need to differentiate between elementary aggregates.[11] The choice for $\chi = 1.25$ means that if, for example, $N_a^{(y,m-1),(y,m)} = 80$, then items with an average expenditure share greater than 1% will be selected.

The price change between *y*,*m*-1 and *y*,*m* for elementary aggregate *a* is now computed as

$$\pi_a^{y,m/y,m-1} = \prod_{i=1}^{n_a^{(y,m-1),(y,m)}} \left(\frac{p_{i,a}^{y,m}}{p_{i,a}^{y,m-1}} \right)^{1/n_a^{(y,m-1),(y,m)}} . \qquad (1)$$

[10] Because the elementary aggregates are chain specific, an item can appear in more than one elementary aggregate with chain-specific prices and expenditure shares.

[11] There is also a practical aspect to it: differentiating would imply determining χ values for several hundreds of elementary aggregates.

Expression (1) is a sample-based month-to-month Jevons index. These monthly indexes are subsequently multiplied (chained) to obtain a long-term time series with an arbitrary reference or starting month y_0, m_0:

$$P_a^{y,m/y_0,m_0} = P_a^{y,m-1/y_0,m_0} \pi_a^{y,m/y,m-1}, \tag{2}$$

where $P_a^{y,m-1/y_0,m_0}$ is the chained matched-items price index going from the starting month to month m-1 of year y. For items that are not available in month y,m but which were purchased in previous periods, a price is imputed as follows:[12]

$$\hat{p}_{i,a}^{y,m} = p_{i,a}^{y,m-1} \pi_a^{y,m/y,m-1}. \tag{3}$$

For higher aggregates A short-term price indexes are calculated according to a fixed-weight ("Laspeyres-type") formula with index reference period y-1:

$$P_A^{y,m/y-1} = \frac{\sum_{a \in A} w_a^{y-1} P_a^{y,m/y-1}}{\sum_{a \in A} w_a^{y-1}}. \tag{4}$$

where

$$P_a^{y,m/y-1} = \frac{P_a^{y,m/y_0,m_0}}{\frac{1}{12}\sum_{s=1}^{12} P_a^{y-1,s/y_0,m_0}}. \tag{5}$$

The weights w_a^{y-1} in (4) are based on the annual expenditures of all items belonging to elementary aggregate a, regardless whether items were included in the sample or not.[13]

Next, the short-term series are chained in December, which is the link month, to construct long-term series with index reference period 0.[14]

$$P_{ch,A}^{y,m/0} = \left(\frac{P_A^{y,m/y-1}}{P_A^{y-1,12/y-1}} \right) \left[\prod_{\tau=1}^{y-1} \frac{P_A^{\tau,12/\tau-1}}{P_A^{\tau-1,12/\tau-1}} \right] P_A^{0,12/0}. \tag{6}$$

Short-term indexes $P_A^{y,m/y-1}$ and chained indexes $P_{ch,A}^{y,m/0}$ are computed at all COICOP levels, both for each supermarket chain separately and across all chains, using formulae (4) and (6).

[12] The previous price (in month y,m-1) can also be an imputed price.

[13] At the 5-digit COICOP level, monthly varying weights are used for seasonal items, such as fresh fruit and fresh vegetables. At the 4-digit COICOP level, fixed annual weights are used.

[14] Currently the index reference period of the Dutch CPI is 2006.

3. A Brief Overview of the RYGEKS Method

The method described in section 2 was chosen by Statistics Netherlands to circumvent the problem of drift in chained weighted (preferably superlative) indexes. An alternative solution, which does make use of a superlative index number formula, was put forward by Ivancic, Diewert and Fox (2009). They adapted the GEKS procedure, known from the international (spatial) price comparisons literature, to price comparisons across time. In this section we will briefly outline their method.

The GEKS method takes the geometric mean of the ratios of all bilateral indexes (calculated using the same index number formula) between a number of entities, such as countries, where each entity serves as the base. Let P^{jl} and P^{kl} be the bilateral indexes between a pair of entities *j* and *l* $(l = 1,...,M)$ and between entities *k* and *l*, respectively. The GEKS index between *j* and *k* can be written as

$$P_{GEKS}^{jk} = \prod_{l=1}^{M} \left[P^{jl} / P^{kl} \right]^{1/M} = \prod_{l=1}^{M} \left[P^{jl} \times P^{lk} \right]^{1/M} . \qquad (7)$$

The second expression of (7) holds when the bilateral price indexes satisfy the "entity reversal test", so that $P^{kl} = 1/P^{lk}$. It can easily be shown that

$$P_{GEKS}^{jk} = P_{GEKS}^{jl} / P_{GEKS}^{kl} . \qquad (8)$$

Expression (8) says that the GEKS price index satisfies the circularity or *transitivity* requirement: the same result is obtained if entities are compared with each other directly or via their relationships with other entities.

For price comparisons across time, the entities in (7) are time periods, for CPI purposes usually months. Consider a range [0,*T*] of $T+1$ months, where *T* is the current (most recent) month. The GEKS index for month *t* $(t \leq T)$ with reference period 0 then is

$$P_{GEKS}^{0,t} = \prod_{\tau=0}^{T} \left[P^{0,\tau} / P^{t,\tau} \right]^{1/(T+1)} = \prod_{\tau=0}^{T} \left[P^{0,\tau} \times P^{\tau,t} \right]^{1/(T+1)} . \qquad (9)$$

One month later *T*+1 is the current month and the time series is calculated for [0,*T*+1]. The GEKS index for month *t* $(t \leq T+1)$ with reference period 0 now becomes:

$$P_{GEKS}^{0,t} = \prod_{\tau=0}^{T+1} \left[P^{0,\tau} / P^{t,\tau} \right]^{1/(T+2)} = \prod_{\tau=0}^{T+1} \left[P^{0,\tau} \times P^{\tau,t} \right]^{1/(T+2)} . \qquad (10)$$

A problem with this approach is that the results for all months $1,...,T$ calculated using (10) will differ from those calculated using (9). So when the observation period is extended and new data is added, the GEKS index numbers will be subject to continuous *revision*, something which is unacceptable for a CPI. A solution to this problem would be to use the index change between T and T+1, calculated using (10), to update the time series (9) through chain linking. By construction, this monthly GEKS price change is affected by the prices and quantities of all months $0,...,T+1$, hence including months $0,...,T-1$. If we would continue to extend the time series in this way, the index change between the last two months will be influenced by the prices and quantities pertaining to more and more months in the past. This involves a loss of "characteristicity". Ivancic, Diewert and Fox (2009) therefore suggested a *rolling year* approach.

The rolling year approach uses a 13 month moving window to compute GEKS indexes. A 13 month window was chosen because it is the shortest window that allows a comparison of strongly seasonal items. The GEKS index (9) for periods [0,12] remains the starting point. As above, the next month-to-month index change is chain linked to the GEKS time series. This index change is computed as the ratio of the GEKS indexes, based on data of [1,13], for months 13 and 12 with reference period 1, which are given by

$$P_{GEKS}^{1,13} = \prod_{\tau=1}^{13} \left[P^{1,\tau} / P^{13,\tau} \right]^{1/13}; \tag{11}$$

$$P_{GEKS}^{1,12} = \prod_{\tau=1}^{13} \left[P^{1,\tau} / P^{12,\tau} \right]^{1/13}. \tag{12}$$

Because GEKS indexes are transitive, the following relations hold:

$$P_{GEKS}^{12,13} = P_{GEKS}^{12,1} / P_{GEKS}^{13,1} = P_{GEKS}^{1,13} / P_{GEKS}^{1,12}. \tag{13}$$

This means that, instead of taking the ratio of (11) and (12), the GEKS index between months 12 and 13, i.e.

$$P_{GEKS}^{12,13} = \prod_{\tau=1}^{13} \left[P^{12,\tau} / P^{13,\tau} \right]^{1/13} = \prod_{\tau=1}^{13} \left[P^{\tau,13} / P^{\tau,12} \right]^{1/13} \tag{14}$$

can be used to update the existing time series. The expression for the rolling year GEKS (RYGEKS) index pertaining to month 13 then becomes

$$P^{0,13}_{RYGEKS} = P^{0,12}_{GEKS} \prod_{\tau=1}^{13} \left[P^{12,\tau} / P^{13,\tau} \right]^{1/13} = \prod_{\tau=0}^{12} \left[P^{0,\tau} / P^{12,\tau} \right]^{1/13} \prod_{\tau=1}^{13} \left[P^{12,\tau} / P^{13,\tau} \right]^{1/13} . \qquad (15)$$

The general expression for the RYGEKS index pertaining to the current month T $(T > 12)$ with reference month 0 is[15]

$$P^{0,T}_{RYGEKS} = \prod_{\tau=0}^{12} \left[P^{0,\tau} / P^{12,\tau} \right]^{1/13} \prod_{t=13}^{T} \prod_{\tau=t-12}^{t} \left[P^{t-1,\tau} / P^{t,\tau} \right]^{1/13} , \qquad (16)$$

and the general expression for the monthly index change is

$$\prod_{\tau=t-12}^{t} \left[P^{t-1,\tau} / P^{t,\tau} \right]^{1/13} = \left[\frac{P^{t-12,t}}{P^{t-12,t-1}} \times \frac{P^{t-11,t}}{P^{t-11,t-1}} \times ... \times \frac{P^{t-3,t}}{P^{t-3,t-1}} \times \frac{P^{t-2,t}}{P^{t-2,t-1}} \times \left(P^{t-1,t} \right)^2 \right]^{1/13} . \qquad (17)$$

Formula (17) provides useful insight into the actual calculation procedure. The monthly change of the RYGEKS index simply equals the (geometric) mean of the direct index between t-1 and t, $P^{t-1,t}$, which is counted twice, and 11 indirect indexes between these months.

As mentioned before, the price indexes that are inputs in the (RY)GEKS method should preferably be superlative indexes. These indexes are symmetric, are grounded in microeconomic theory and have good axiomatic properties, including the time reversal requirement. Ivancic, Diewert and Fox (2009) used Fisher price indexes. In a follow-up study, de Haan and van der Grient (2009) used Törnqvist indexes instead. The recently built shadow scanner data system at Statistics Netherlands is also based on Törnqvist-type RYGEKS indexes.

4. Scanner Data and Data Cleaning

4.1 Data Description

At present, eight supermarket chains send data files to Statistics Netherlands every week containing scanner data on all individual items sold, seven of which are currently used in the CPI. Items are identified by the European Article Number (EAN). Each record of a data file relates to a single EAN and contains weekly expenditures and quantities sold. It also contains a short product description, often including pack size and package size. Records that have identical EANs are considered to refer to the exact same (physical)

[15] In our empirical illustration, formula (16) will be applied at the 5-digit-COICOP level.

item. Items with different EANs are treated as different products. The consequences of these choices are addressed in section 6.

Each chain adds its own classification code indicating to which product category the EANs belong. Having a classification code attached to the data is indispensable for an efficient CPI process given the overall huge attrition rate of EANs. Once the relation between the chain-specific classification and COICOP has been established, EANs are automatically assigned to the appropriate 5-digit COICOP category. To prevent regular recoding, a prerequisite is that the chain-specific classifications are stable through time. Of course they should also be more detailed than the 5-digit COICOP level.

Supermarket chains sometimes group together items related to special occasions like children's birthdays. Products that serve a particular aim, such as meat, charcoal and sauces bought when organising a barbecue, may also be grouped together. In these cases the EANs cannot automatically be assigned to one of the COICOP categories. For efficiency reasons it was decided not to assign the items to the appropriate category but to simply exclude them.

Supermarkets considerably expanded the range of products sold during the past decades. The assortment presently includes, for instance, clothing, glassware, tableware and household utensils. For the time being Statistics Netherlands decided to restrict the calculation of scanner data indexes to the more traditional product categories. Table 1 lists the COICOP categories for which scanner data indexes are now calculated.

Table 1. COICOP categories for which scanner data indexes are calculated

Code	Description
01	Food and non-alcoholic beverages
02.1.2	Wine
02.1.3	Beer
05.5	Tools and equipment for house and garden
05.6	Goods and services for routine household maintenance
06.1	Not reimbursable medical and pharmaceutical products
09.3.4	Pets, pet foods and products for pets
13.1.3	Appliances, articles and products for personal care

4.2 Data Cleaning and Data Preparation

Timeliness is an important aspect of a CPI. The Dutch CPI for month *t* is published in the first week of month *t*+1. A consequence is that data referring to the last week of month *t* cannot be used. Price collection in the field is therefore restricted to the first three full weeks of each month. A similar procedure is applied to scanner data. That is, the unit value of an item is based on the expenditures and quantities sold pertaining to the first three full weeks of a calendar month from all stores within a chain for which scanner data is received. The same goes for the average expenditure shares used in the (cut-off) sampling procedure.

The prices (unit values) are subject to two automatic data cleaning procedures. First, month-to-month price relatives greater than 4 or smaller than ¼ are considered implausible and declared invalid. Thus, items for which the current price is more than 300% higher or more than 75% lower than the previous price are automatically deleted. Second, an algorithm was developed, referred to as a dumping filter, that excludes items from the computation which exhibit strong price decreases in combination with strong decreases in expenditures. “Dumping” occasionally occurs in case of stock clearances when an item is sold at an extraordinary low price. Since the item will not be available any longer, it does not return to a regular price. The resulting price decreases, without offsetting price increases, can have an undesirable downward effect on the index of the product category in question. In practice this dumping filter may also delete some items that are not related to stock clearances. This should not be viewed as a serious problem as missing prices are imputed.

While it can be argued that the “dumping filter” should be augmented for the computation of RYGEKS indexes, or that it should not be used at all, we implemented it in the shadow scanner data system without making any changes. This makes it possible in section 5 below, where we compare the results of the two methods, to focus solely on the impact of the choice of index number formula.

5. A Comparison of the Two Methods

5.1 The 5-digit COICOP Level

The shadow system computes RYGEKS indexes for 58 different product categories at the 5-digit COICOP level, being the lowest publication level in the Dutch CPI (though the scanner data indexes are not published separately). At this level, price indexes are computed for each of the six supermarket chains that are in the Dutch CPI since January 2010. However, some of the chains do not sell products of all 58 categories, so that 343 cross classifications of product groups and supermarket chains remain, also referred to as strata. Note again that 5-digit COICOP is one level higher than the elementary "6-digit COICOP" level for which Jevons indexes in the Dutch method are compiled. We have chosen the 5-digit level to increase the number of observations. Since RYGEKS indexes are based on (bilateral) superlative index numbers, there seems to be no need to compute them at the lower aggregation level.

In the present study we compare monthly RYGEKS and "official" price index numbers, covering the period January 2009, which serves as the index reference period, to October 2010. The price index numbers for the 343 different strata using the Dutch method were extracted from the official CPI database and re-scaled to make them equal to 100 in January 2009. For each of the 343 strata we computed the average difference between the RYGEKS and the "official" index numbers during the 22-month period. Table 2 provides a summary of the findings. For 81% of the strata (278 out of 343) the average difference is less than 2 index points, for 62% even less than 1 index point. If we take RYGEKS as the benchmark, we may be tempted to conclude that the Dutch method performs reasonably well.

Table 2. Average difference between official and RYGEKS indexes at the 5-digit COICOP level

Average difference	Number of strata	
	absolute	%
(-14, -4)	14	4
[-4, -3)	5	2
[-3, -2)	14	4
[-2, -1)	21	6
[-1, 1]	211	62
(1, 2]	46	13
(2, 3]	15	4
(3, 4]	11	3
(4, 12]	6	2

On the other hand, there are 20 strata where the average difference of the index numbers according to the two methods exceeds 4 index points. Differences that large cannot be ignored. Table 3 indicates that 9 out of these 20 strata pertain to the 4-digit COICOP categories "fruit" and "vegetables and potatoes", product categories which are dominated by fresh produce. In a few instances the average absolute differences even exceed 10 index points.

Table 3. COICOP (4-digit) categories where official and RYGEKS stratum indexes differ more than 4 index points on average

COICOP	Description	Number of strata
01.1.2	Meat	2
01.1.3	Fish	2
01.1.6	Fruit	3
01.1.7	Vegetables and potatoes	6
01.2.2	Mineral waters, soft drinks, fruit and vegetable juices	1
02.1	Alcoholic beverages	3
05.6.1	Non-durable household goods	2
12.1.3	Articles and products for personal care	1

5.2 Higher Aggregates

The RYGEKS indexes for the 343 (product category x supermarket chain) strata have been aggregated to higher COICOP levels using fixed weights that relate to the year 2009. The same procedure was used to aggregate the "official" indexes. It should be noted that this upper-level weighting procedure differs from the procedure applied in the Dutch CPI where, as was described in section 2, annual chaining is used. We applied this aggregation method to the "official" stratum indexes also, so that the aggregation method does not affect the comparisons.[16]

Figure 1 shows the all items, all supermarket chains scanner data price indexes, i.e. the price indexes aggregated across all strata, according to both methods. There is no clear sign of a persistent difference. The RYGEKS index is on average 0.08 points above the index based on the (approximate) Dutch method. Apparently the differences at the stratum level cancel out. Thus if the RYGEKS approach is considered ideal, the Dutch method performs satisfactorily at the highest level of aggregation.

As can be seen from Figure 1, supermarket prices in the Netherlands slightly went down during 2009-2010. Although the observation period is too short to draw any definitive conclusions, there seems to be a seasonal pattern with prices being relatively high in spring and relatively low in autumn. The most striking feature of Figure 1 is the

[16] We chose this aggregation procedure for simplicity. It gives rise to a Young-type index, which can be problematic. A better choice would have been to "price backdate" the expenditure shares to January 2009, leading to a Lowe-type index.

sharp drop in prices in January 2010. As we found in earlier years as well, January is a popular month for promotional sales.

Figure 1. Aggregate price change (58 product categories, 6 supermarket chains)

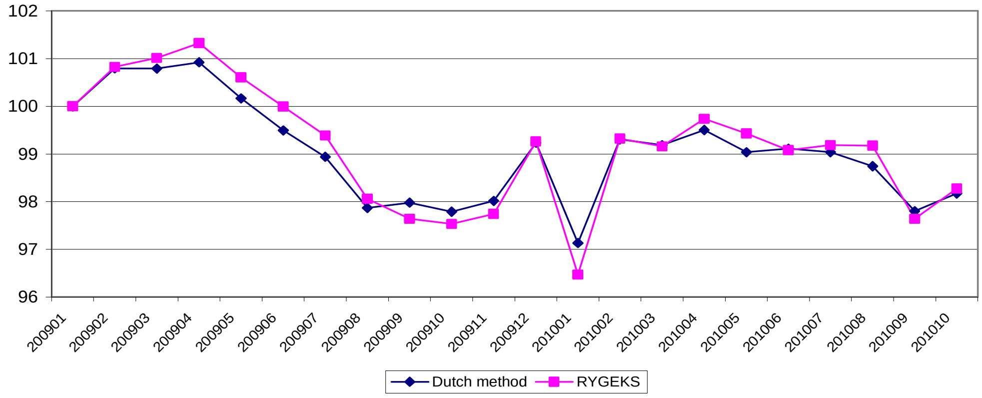

The impact of sale prices in January 2010 can also be observed for many product categories at the 4-digit COICOP level; see Figure 2. Not surprisingly, measured price changes at this level are more volatile than the price change at the most aggregate level. For some COICOP classes, for example "meat" and "food products n.e.c.", the month-to-month fluctuations are substantial. In most cases the impact of the method used on the 4-digit results is limited to a few index points but there are some notable exceptions. For "fruit" the index numbers using the Dutch method differ up to 4 points from those using RYGEKS. Yet, the trends are similar and the differences seem to be incidental rather than structural. In several other cases, e.g., for "alcoholic beverages" and "non-durable household goods", the differences are persistent and the Dutch method clearly understates the RYGEKS price development. As will be discussed later on, while those differences are persistent, they do not always grow over time but may be caused by one-off incidents.

Figure 2. Scanner data price indexes indexes; COICOP 4-digit level

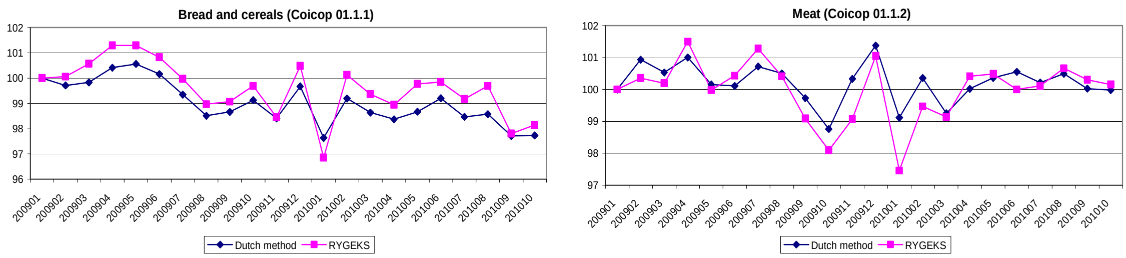

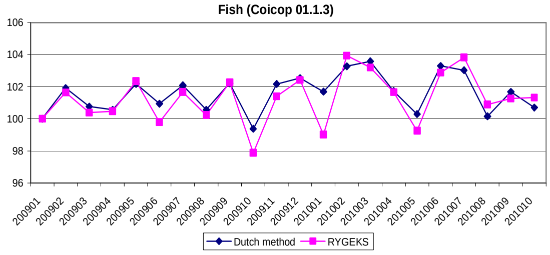
Fish (Coicop 01.1.3)
106
104
102
100
98
96
Dutch method
RYGEKS

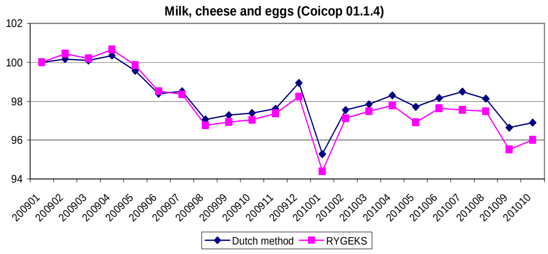
Milk, cheese and eggs (Coicop 01.1.4)
102
100
98
96
94
Dutch method
RYGEKS

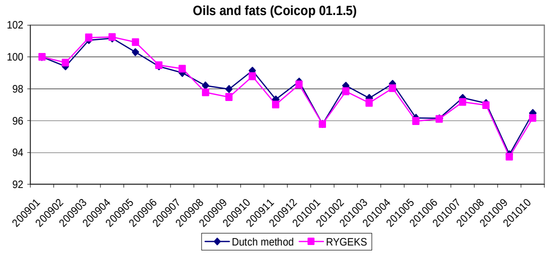
Oils and fats (Coicop 01.1.5)
102
100
98
96
94
92
Dutch method
RYGEKS

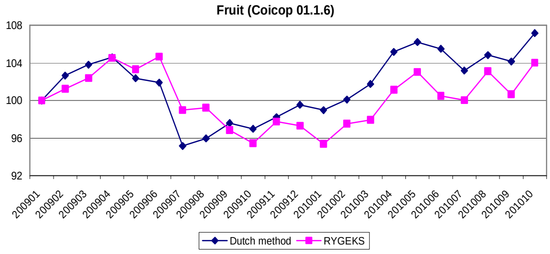
Fruit (Coicop 01.1.6)
108
104
100
96
92
Dutch method
RYGEKS

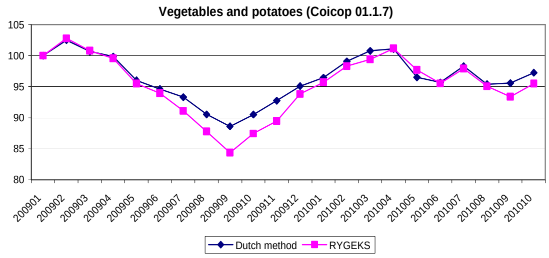
Vegetables and potatoes (Coicop 01.1.7)
105
100
95
90
85
80
Dutch method
RYGEKS

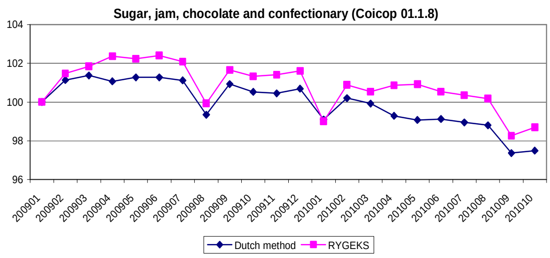
Sugar, jam, chocolate and confectionary (Coicop 01.1.8)
104
102
100
98
96
Dutch method
RYGEKS

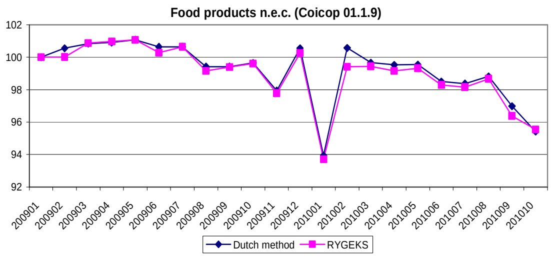
Food products n.e.c. (Coicop 01.1.9)
102
100
98
96
94
92
Dutch method
RYGEKS

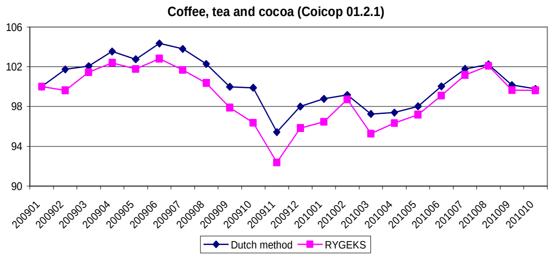
Coffee, tea and cocoa (Coicop 01.2.1)
106
102
98
94
90
Dutch method
RYGEKS

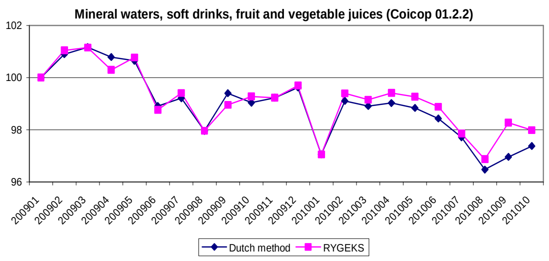
Mineral waters, soft drinks, fruit and vegetable juices (Coicop 01.2.2)
102
100
98
96
Dutch method
RYGEKS

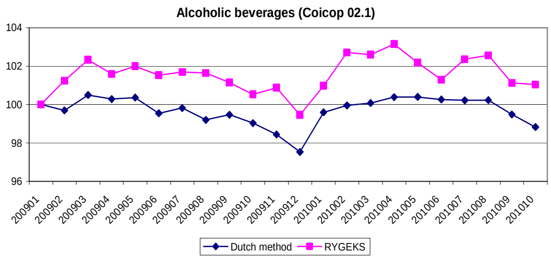
Alcoholic beverages (Coicop 02.1)
104
102
100
98
96
Dutch method
RYGEKS

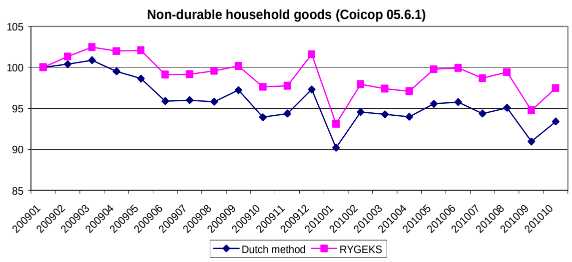
Non-durable household goods (Coicop 05.6.1)
105
100
95
90
85
Dutch method
RYGEKS

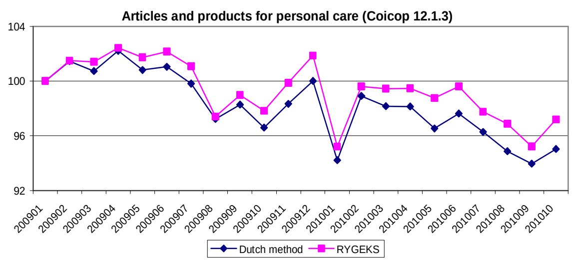
Articles and products for personal care (Coicop 12.1.3)
104
100
96
92
Dutch method
RYGEKS

The "RYGEKS" indexes in Figures 1 and 2 have been calculated as weighted averages of (5-digit COICOP) RYGEKS indexes, using a fixed-weight index number formula. That is, the price indexes at all levels above 5-digit COICOP are in fact hybrid indexes. The advantage of this two-stage procedure is that, at the upper level, the price indexes are consistent in aggregation.[17] However, the RYGEKS methodology can be applied directly to any aggregation level. As a matter of fact, from a theoretical point of view it would be better to compute RYGEKS indexes at all levels because this would not only account for substitution within product categories but also across categories.

To get an idea of what the impact might be of keeping upper level weights fixed, we computed for each supermarket chain indexes for COICOP division 01 ("food and non-alcoholic beverages"), using the RYGEKS approach directly as well as using the hybrid versions where RYGEKS was applied at the 5-, 4- and 3-digit level. All items, all chains indexes were calculated using the fixed-weight formula. Figure 3 shows the differences between these hybrid versions and the "true" RYGEKS index.

Figure 3. The effect of using RYGEKS at different levels of aggregation

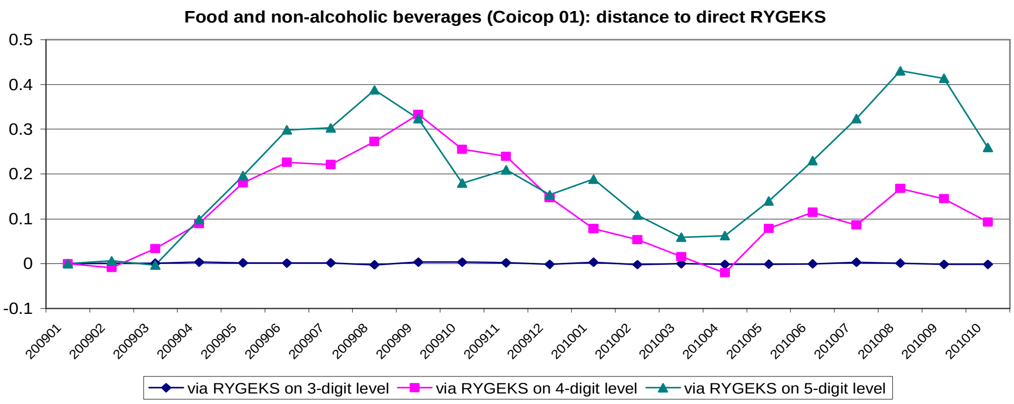

Keeping weights fixed at the 3-digit level only has a negligible effect. But when the weights are kept fixed at lower levels also, a seasonal pattern emerges that is mainly due to "fruit and vegetables". Expenditures on these seasonal goods are much greater in summer than in winter, and the prices of fresh produce are lower in summer, which has a downward effect on the "true" RYGEKS index.

[17] This is true because we are using fixed weights. If we had used annual chaining, as in the official Dutch method, then the long-term (chained) indexes would not have been consistent in aggregation.

6. Issues When Dealing with Scanner Data

(Lack of) Weighting

Although the Dutch method and the RYGEKS method differ in other respects as well, by far the most important reason why they generate different results, especially at lower aggregation levels, is the treatment of the relative importance of the items. The lack of expenditure weights in the Dutch method for the items included in the cut-off samples at the lowest (elementary) aggregation level is of course the biggest issue. This becomes most apparent for elementary aggregates where price changes differ substantially among items and where the expenditure distribution is highly skewed.

Fresh fruit is a typical example. A large proportion of the expenditures on fresh fruit in the Netherlands is spent on bananas; hence the impact of bananas will be too low in the Dutch method. During the first half of 2009 and the first half of 2010, the price of bananas increased a lot less than the price of other fresh fruit. This explains to a large extent the difference in price development according to both methods shown in Figure 4, which relates to a single supermarket chain ("C"). The index based on the Dutch method clearly has an upward bias.

Figure 4. Fresh fruit (COICOP 01.1.6.1), supermarket chain "C"

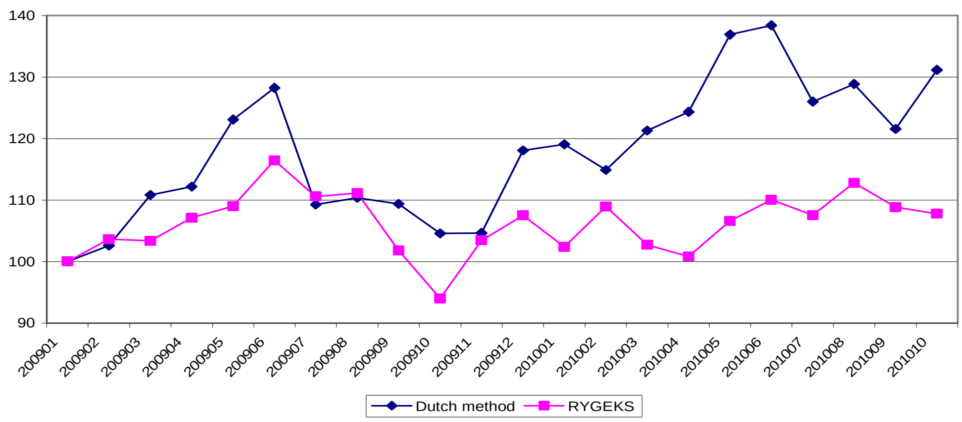

The Use of Price Filters

As was mentioned in section 4, two data cleaning procedures, or "price filters", are used in the Dutch method and have also been implemented in the shadow RYGEKS system, a "dumping filter" and a filter that excludes items with implausibly large price changes. The usefulness of the filters is illustrated in Figures 5 and 6. These figures show the all items indexes (pertaining to 58 5-digit COICOP categories) for two supermarket chains "E" and "B", respectively, based on the Dutch method (including the two price filters) and based on the RYGEKS method, both with and without using filters. The raw data for chain "E" registered a huge price rise for a particular item between April and May 2009. Not applying price filters would have caused a RYGEKS increase of 3% at the all items level, as can be seen in Figure 5. The data entry error – which we think it was – was a one-off incident as in June 2009 the index returns to a "normal" level. The price filters effectively remove the implausible spike.

Figure 5. All items price indexes; supermarket chain "E"

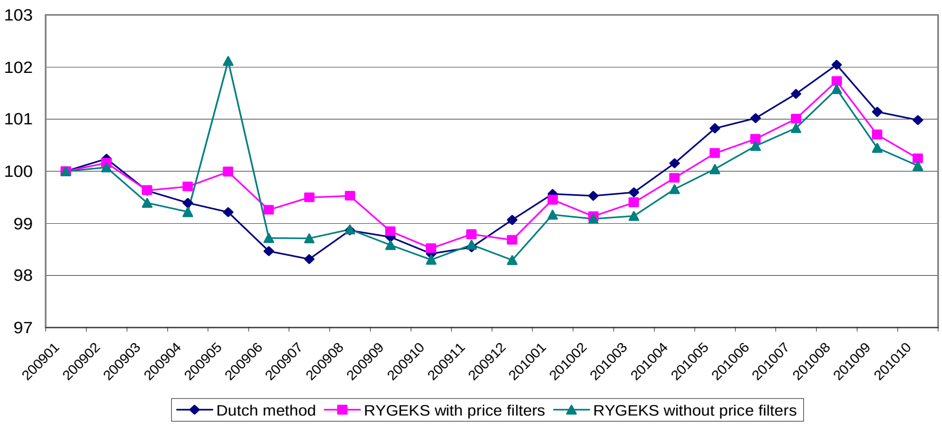

While the data entry error in the example of Figure 5 would not have affected the trend, unnoticed changes in the way the data is registered can also cause structural breaks. An example is given in Figure 6. In May 2010, supermarket chain "B" suddenly changed the unit of measurement for many items from a price per (half) kilo to a price per gram. Most of these items were given a new EAN but for a limited number of items the EAN was left unchanged. Since the EAN is the item identifier, the latter items were matched and showed an incorrect price decrease by a factor of 500 or 1000. Without the use of the price filters, the RYGEKS method would have measured a price decrease of 15% at the all items level for chain "B", as shown in Figure 6. Again, the filters effectively remove the implausible price change.[18]

[18] Obviously, it would have been better if we had repaired the error by re-computing the prices in constant units of measurement because the number of matched items now is much lower than the "true" number. However, that would have been a lot of work. See also section 6.4.

Figure 6. All items price indexes; supermarket chain "B"

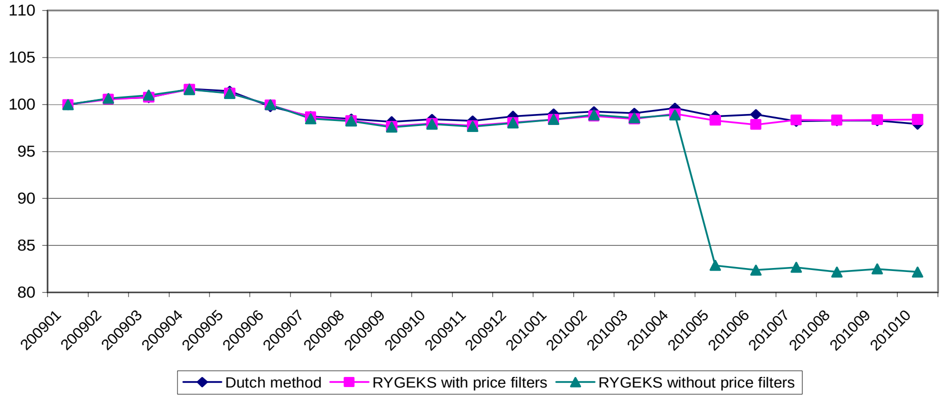

Imputations for Temporarily "Missing Prices"

A feature of the Dutch method is that the "missing prices" for temporarily unavailable items are imputed. When an item has disappeared and re-appears after a while, the new price is compared with an imputed previous price, ensuring that in the longer term the true price change is captured. Without imputations the price change during the period of absence would be missed. Since the RYGEKS method takes into account price changes of matched items during a whole year to compute the most recent monthly price change, it may not be necessary to impute prices explicitly.[19]

The EAN as Item Identifier

Records identified by the same EAN are considered to refer to the exact same (physical) item. EANs are assigned to specific products by the manufactures, and for the majority of the items sold in supermarkets we can be confident that a particular EAN in different stores and in different time periods relates to a single homogeneous product, at least during a couple of years. For some products, such as fresh fruit, supermarkets assign unofficial "EANs" themselves, either at the store level or at the chain level. A specific set of "EANs" is available for this purpose; shorthand codes are sometimes used at the checkout. The use of these store-specific or chain-specific "EANs" does not pose any problems for price index construction as long as they refer to the same items over time. However, identical store-specific codes are irregularly assigned to different products in

[19] This is a bit speculative though. We are currently estimating RYGEKS indexes in which the "missing prices" are actually imputed to investigate if the effect of imputations is indeed negligible.

different months, in which case we would not be comparing like with like. Fortunately, this phenomenon happens so rarely, and the impact on the results appears to be so small, that we can safely ignore the problem. Moreover, if the resulting price changes were substantial, they would most likely be eliminated by the price filters in the data cleaning procedures or during the monthly routine of checking and analysing the results of the index computation.[20]

In some instances the EAN level could be too detailed for CPI purposes. Items with different EANs, but which are identical from the consumers' point of view, should in principle be treated as the same product. If an item (identified by its EAN) disappears and a completely comparable item with a different EAN appears, then the prices should be directly compared. An example would be a pack of coffee that is normally wrapped in red paper but, for promotional reasons, has suddenly been wrapped in black paper. Neither the Dutch method nor the RYGEKS method would automatically compare the prices. As mentioned before, the possibility to make explicit adjustments is built into the current computer system. So if the CPI statisticians were to know those situations, they could in principle make direct comparisons.[21]

The problem shown in Figure 6 is an example where new EANs were suddenly assigned to existing items, in this case because the unit of measurement was changed. The best option would obviously be to compare the prices directly after adjusting for the change in the unit of measurement. In reality, the Dutch index was completely based on the (matched) items whose EANs had not changed.

Promotional Sales and Discounts

Supermarkets frequently sell items at prices far below the regular price; discounts can be up to 50%. Consumers typically react to the relative price changes by purchasing the items that are on sale in large quantities, in particular when the items can be stored at home. The RYGEKS method was in fact designed to cope with this situation. Because the increase in expenditures and the decrease in a later period are explicitly taken into account, the RYGEKS elementary indexes often turn out to be more volatile than the

[20] For some supermarkets we know which "EANs" are used for store- or chain-specific coding, and we might decide to exclude them from the index calculation in the future.

[21] The computer system provides the user with indicators pointing to major changes in the assortment of the supermarket chain in question, which might be helpful in this respect.

corresponding indexes according to the unweighted Dutch method. Figure 7 provides a nice illustration. Supermarket chain “C” offered a particular brand of rice at discount prices in January, April and September 2010. As a result the RYGEKS index behaves in an erratic way. This volatility is not necessarily undesirable. If the aim is to estimate an acquisitions index (like in Australia, for example), as opposed to a cost-of-use or cost-of-living index, it can be argued that we ought to measure the instantaneous impact of the purchases made.

Figure 7. Volatility of the RYGEKS index; Rice (COICOP 01.1.1.1), supermarket chain “C”

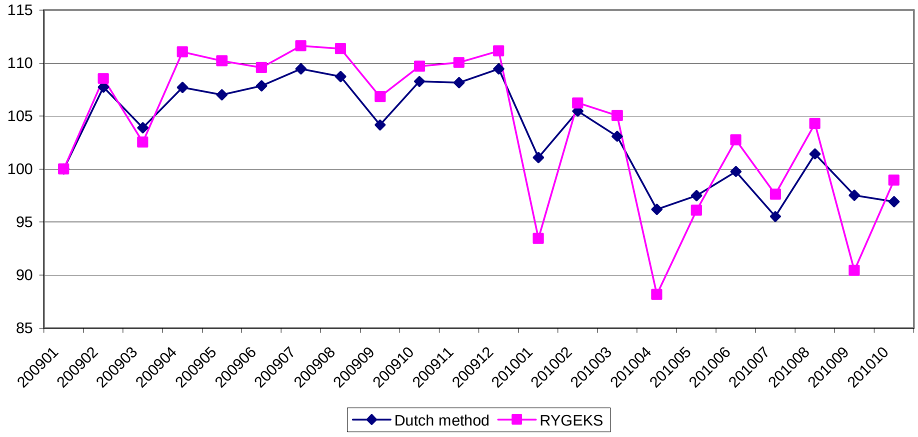

Promotional prices can cause a *starting problem*, as we call it, in the RYGEKS index. This phenomenon occurs when the month of introduction of a supermarket chain into the index calculation includes a sales period (which in the Netherlands usually endures a full calendar week). During the month after the introduction, the prices tend to return to regular levels, and the RYGEKS index will show a strong increase. Obviously, an unweighted index will be much less affected. This is clearly demonstrated in Figure 8 for the price change of “wine” in supermarket chain “A”. Because the price decrease between December 2008 and January 2009 is unobserved, the RYGEKS index has an “upward bias” during the whole observation period.

Of course the problem can simply be avoided by incorporating the chain, or the elementary aggregate(s) in question, one month later, but practical considerations may hamper this. The starting problem is a general problem with linking in volatile index series. The reason we have discussed it is that one should be aware of the problem when interpreting the results and comparing the two methods.

Figure 8. The starting problem; Wine (COICOP 02.1.2), supermarket chain "A"

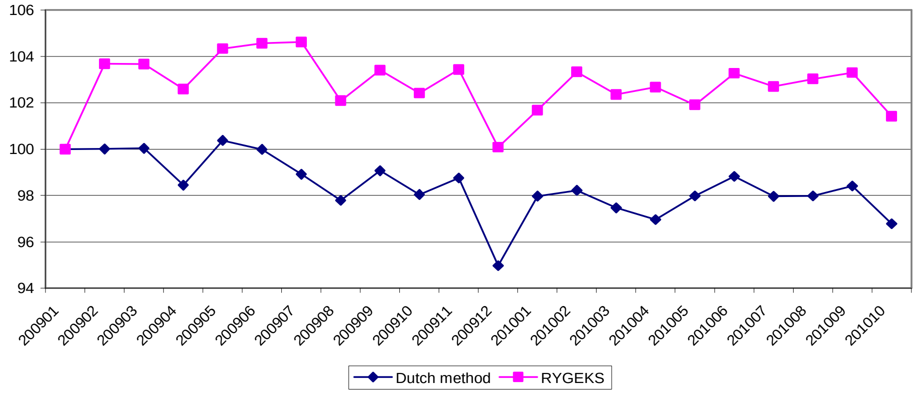

Large price cuts do not only happen in case of temporary promotional discounts, they can also be structural. Figure 9 shows an example for "milk" sold in supermarket chain "F". In May and June 2009 chain "F" lowered the prices of several items. This generated a large increase in turnover, presumably by attracting many new customers as milk is not a storable product. The price cuts and the resulting increase in expenditures had a downward effect on the RYGEKS index of almost 10%. The (unweighted) Dutch method measures a much smaller price decrease. The difference between both indexes remains more or less constant over time because prices did not return to their original levels. Figure 9 once again illustrates the problem of using an unweighted index number formula.

Figure 9. The effect of structural price cuts; Milk (COICOP 01.1.4.1), supermarket chain "F"

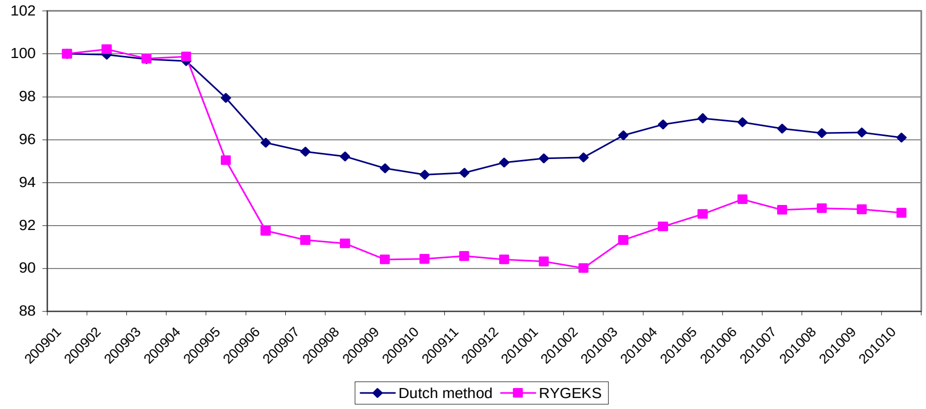

Strongly Seasonal Items

Strongly seasonal items are items which are only available during certain periods of the year. Fresh fruit like strawberries and vegetables are typically mentioned, but there are more items that have a strong seasonal pattern. Prior to Easter, Christmas and a specific Dutch December event ("Sinterklaas") consumers purchase large quantities of chocolate delicacies. These products are on the market for two or three consecutive months only. After the events, i.e., in the second or third month, the last remains are sold at clearance prices. A year later the items may re-appear on the market, and most of them appear to have the same EAN as last year so that prices can be directly compared.

Figure 10 shows the price change of the pseudo 6-digit COICOP product group "Chocolate" sold by supermarket chain "F". In the relevant months, like November and December, the seasonal delicacies dominate the expenditures. The product category also contains items that are available throughout the year. Given that the RYGEKS method explicitly compares current prices with those of the previous year, we would expect the method to work fine.[22] In any case, the RYGEKS indexes are very stable through time, which is reassuring since at the item level we did not observe major price changes during twelve-month periods, neither for the seasonal items nor for the standard ones. The Dutch method compares the prices of the seasonal items implicitly by means of imputations. While measured price changes are similar for both methods during sub-periods, by the end of the period the RYGEKS index is 10 points higher.

Figure 10. Strongly seasonal items; Chocolate (COICOP 01.1.8.3.1), supermarket chain "F"

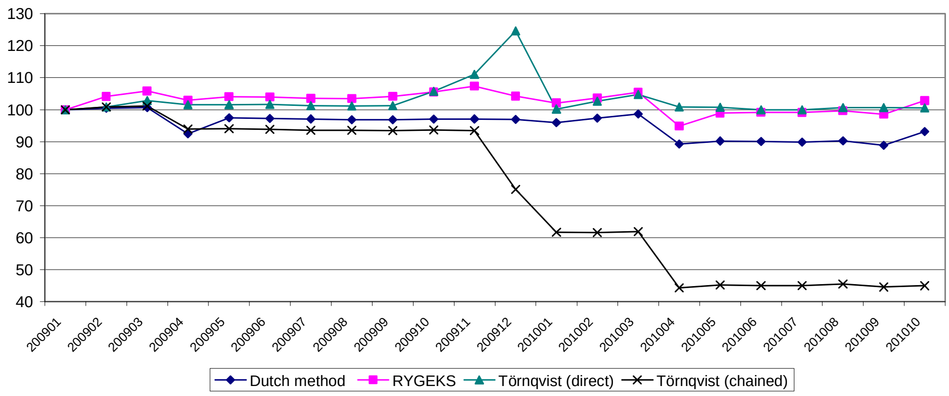

[22] Balk (1981) proposed the use of the GEKS method for seasonal goods in a producer price index context already thirty years ago, though he did not propose a rolling year approach.

For this example we also computed direct and monthly chained Törnqvist price indexes (without imputations). As can be seen from Figure 10, the chained Törnqvist has a severe downward bias. That chaining can lead to drift when there is seasonality in the prices is well known. In this case the prices of the seasonal delicacies systematically fall during adjacent months and these price decreases are heavily weighted. Actually, this example illustrates that, without imputations, monthly matching and chaining does not capture the implicit (but unobserved) offsetting price increases. The direct Törnqvist index compares the prices of items purchased in the current month and in the starting period, which is January 2009. Thus, this index does not take the Easter delicacies into account. The strong increase during the last months of 2009 stems from the fact that the regular prices of Christmas 2009 delicacies are compared with the low clearance prices in January 2009.

7. Assessment of the Two Methods

Our empirical comparison of the Dutch method for treating supermarket scanner data with the alternative RYGEKS method was meant in the first place to provide input to Statistics Netherlands management to make an informed decision about implementation of the RYGEKS method. Unfortunately the shadow system that currently computes RYGEKS price indexes does not meet official IT requirements. Implementation would therefore mean (re-)starting a costly IT project.

Based on our findings, below we assess the two methods by looking at three aspects: accuracy of the method, technicalities and consequences for the CPI production process.

Accuracy

From a theoretical point of view the RYGEKS method is undoubtedly better than the Dutch method. The RYGEKS method produces weighted indexes, is based on bilateral superlative indexes (which are grounded in economic theory and have good axiomatic properties), and makes optimal use of all the available price and quantity information without suffering from chain drift. Our empirical evidence supports this view. While at the all items, all supermarket chains level, the Dutch scanner data index is very similar to the RYGEKS, at lower aggregation levels we found some examples where the Dutch method produced rather implausible results. For some product categories, such as fresh fruit, the differences with the corresponding RYGEKS indexes are persistent and grow over time.

The RYGEKS method, as it is a matched-model approach, does not explicitly adjust for quality changes. Whether this gives rise to quality change bias is difficult to say. A matched-model approach does not necessarily lead to bias; much depends on the market circumstances. Given the degree of competition in the Dutch market, it may well be that explicit adjustments would be superfluous. Furthermore, we do not believe that quality improvements for goods purchased in supermarkets are substantial (at least not compared with high-tech goods). Statistics Netherlands decided not to explicitly adjust for quality changes when implementing the current method. We would advice to do the same if the RYGEKS method is to be implemented.

Technicalities

An artificial “6-digit COICOP” aggregation level was introduced in the Dutch method to allow for weighting below the 5-digit COICOP level. This artificial aggregation level is differs across supermarket chains, which is a bit impractical for further aggregation steps. Also, the weighting scheme at the “6-digit” level has to be updated every year since annual chaining is used. The RYGEKS method can be applied directly to the 5-digit COICOP level – there is no need for further disaggregation unless users would like to have more detailed figures in the future, which is unlikely – and the resulting price index numbers can be transferred directly to the CPI aggregation module. In addition, the RYGEKS method can and probably should be applied to all items belonging to a 5-digit COICOP category (as we have done) instead of a sample thereof. Thus, the cut-off sampling procedure can be dropped. This, and dropping the artificial “6-digit COICOP” level, would make the system simpler than the current one.

A drawback of using the RYGEKS method might be the increased volatility of the time series and the starting problem this could cause. On the other hand, as was mentioned earlier, an increase in volatility is not undesirable if this describes the real world. The starting problem, though bigger for RYGEKS than for the Dutch method, is a general problem arising from the use of volatile index series for which a remedy can easily be found.

Consequences for the CPI Production Process

Although the system would be simpler in several respects when using the RYGEKS method, the actual index number calculation would be more complex, if only because the data pertaining to 13 instead of 2 months are used to compute the most recent index change. The RYGEKS method is completely different from methods applied so far, and it will probably take quite some time for practitioners to get familiar with the RYGEKS formula and its properties. This should not be underestimated.

A complicating factor might be that it is not possible to decompose the (change in the) RYGEKS index exactly into contributions of the individual items, as can be done for conventional formulae. Price statisticians often find such a decomposition useful in the validation process of the indexes. Van der Grient (2010) showed that, except under very unusual circumstances, items with the biggest weighted price changes between the last two months will dominate the change in the RYGEKS index. So there seems to be no need for changing the monthly process of validation.

Conclusion

Our advice would be to change over to the RYGEKS method since in our opinion the methodological advantages clearly outweigh any practical drawbacks. The method will however not be implemented this year or next year, due to some of the issues mentioned above. Whether it will come to implementation in later years remains to be seen.

References

Balk, B.M. (1981), A Simple Method for Constructing Price Indices for Seasonal Commodities, *Statistische Hefte* 22, 1-8.

Becker-Vermeulen, C. (2006), Recent Developments in the Swiss CPI: Scanner Data, Telecommunications and Health Price Collection, Paper presented at the ninth meeting of the Ottawa Group, 14-16 May 2006, London.

Feenstra, R.C. and M.D. Shapiro (2003), High Frequency Substitution and the Measurement of Price Indexes, pp. 123-146 in R.C. Feenstra and M.D. Shapiro (eds.), *Scanner Data and Price Indexes*, Studies in Income and Wealth, Volume 61. University of Chicago Press, Chicago.

Fox, K.J. and D. Melser (2011), Non-linear Pricing and Price Indexes: Evidence and Implications from Scanner Data, Discussion paper, School of Economics, University of New South Wales, Sydney.

Greenlees, J. and R. McClelland (2010), Superlative and Regression-Based Consumer Price Indexes for Apparel Using U.S. Scanner Data, Paper presented at the Conference of the International Association for Research in Income and Wealth, St. Gallen, Switzerland, August 27, 2010.

van der Grient, H.A. (2010), How to Analyse Changes in Price Indexes when Applying the RYGEKS Formula, Statistics Netherlands, The Hague.

van der Grient, H.A. and J. de Haan (2010), The Use of Supermarket Scanner Data in the Dutch CPI, Paper presented at the Joint ECE/ILO Workshop on Scanner Data, 10 May 2010, Geneva. Available on: www.cbs.nl.

de Haan, J. (2008), Reducing Drift in Chained Superlative Price Indexes for Highly Disaggregated Data, Paper presented at the ninth Economic Measurement Group Workshop, Sydney, 10-12 December 2008.

de Haan, J. and H.A. van der Grient (2009), Eliminating Chain Drift in Price Indexes Based on Scanner Data, Paper presented at the eleventh meeting of the Ottawa Group, 27-29 May 2009, Neuchâtel.

de Haan, J. and H.A. van der Grient (2011), Eliminating Chain drift in Price Indexes Based on Scanner Data, *Journal of Econometrics* (forthcoming).

ILO, IMF, OECD, Eurostat, United Nations, World Bank (2004), *Consumer Price Index Manual: Theory and Practice*, Geneva: ILO Publications.

Ivancic, L. (2010), Submission to the 16th Series Review of the Consumer Price Index. Available at www.abs.gov.au.

Ivancic, L., W.E. Diewert and K.J. Fox (2009), Scanner Data, Time Aggregation and the Construction of Price Indexes, Discussion Paper no. 09-09, Department of Economics, University of British Columbia, Vancouver, Canada.

Ivancic, L. and K.J. Fox (2010), Understanding Price Variation Across Stores and Supermarket Chains: Some Implications for CPI Aggregation Methods, Discussion paper, School of Economics and Centre for Applied Economic Research, University of New South Wales, Sydney.

Ivancic, L., W.E. Diewert and K.J. Fox (2011), Scanner Data, Time Aggregation and the Construction of Price Indexes, *Journal of Econometrics* (forthcoming).

Rodriguez, J. and F. Haraldsen (2006), The use of scanner data in the Norwegian CPI: The "new" index for food and non-alcoholic beverages, *Economic Survey* 4, 21-28.

Schut, C. (ed.) (2002), Gebruik van scannerdata van supermarkten in de consumentenprijsindex, Statistics Netherlands, Voorburg (Dutch only). Available at www.cbs.nl.

Triplett, J.E. Using Scanner Data in Consumer Price Indexes: Some Neglected Conceptual Considerations, in R.C. Feenstra and M.D. Shapiro (eds.), *Scanner Data and Price Indexes*, Studies in Income and Wealth, Volume 61. University of Chicago Press, Chicago.