

This PDF is a selection from an out-of-print volume from the National Bureau of Economic Research

Volume Title: Scanner Data and Price Indexes

Volume Author/Editor: Robert C. Feenstra and Matthew D. Shapiro, editors

Volume Publisher: University of Chicago Press

Volume ISBN: 0-226-23965-9

Volume URL: <http://www.nber.org/books/feen03-1>

Conference Date: September 15-16, 2000

Publication Date: January 2003

Chapter Title: Scanner Indexes for the Consumer Price Index

Chapter Author: J. David Richardson

Chapter URL: <http://www.nber.org/chapters/c9728>

Chapter pages in book: (p. 39 - 66)

Scanner Indexes for the Consumer Price Index

David H. Richardson

2.1 Introduction

The Consumer Price Index (CPI) is the nation's primary measure of the price change of consumer goods and services. To produce the CPI, the Bureau of Labor Statistics (BLS) staff tracks the prices of a sample of consumer items in the various categories of consumer spending in stores and other retail outlets. In some of those categories, virtually all of the items have a manufacturer-supplied Universal Product Code (UPC) printed on the products to be read by scanners (McKaig 1999). Retailers set prices by UPC for efficiency at the checkout and for inventory management. Consequently, the retailers create computerized records by UPC of the prices and number of units they have sold, records that are commonly called *scanner data*. Scanner data are also used in marketing research to track promotions and variations in packaging, product size, and pricing. In a report to the Senate Finance Committee, Boskin et al. (1996) suggested that scanner data could be used in the CPI for additional commodity detail, as have the Conference Board (1999) and de Haan and Opperdoes (1997a,b). Previous BLS work on scanner data includes Bradley et al. (1997) and Reinsdorf (1997).

The data used in the basic indexes of most item strata are collected by BLS field economists, typically by personal visits to the store or other retail outlet. In retail outlets such as supermarkets, the collection is by observa-

David H. Richardson is an economist at the U.S. Bureau of Labor Statistics working in the CPI program.

Many thanks to all of the members of the ScanData team; this paper is really a group effort. Special thanks to Ralph Bradley, Bill Cook, Dennis Fixler, John Greenlees, Bill Hawkes, and Walter Lane for helpful comments, to Lyubov Rozental for the programming assistance that made it all possible, and to Scott Pinkerton for cereal analysis. All remaining errors are of course the author's own.

tion of the prices on the shelves. Each of these prices is called a quote, whether a price was actually collected or whether an unsuccessful attempt was made. The outlets and items are chosen using a probability of selection proportional to sales (PPS) approach. The outlet sample for most strata is based on the point-of-purchase survey (POPS), a telephone inquiry of where items are purchased by consumers in the urban United States. The result is that the CPI tracks prices in the outlets where people actually shop. Within the sample outlets, the item samples are selected in proportion to the outlets' sales within the item category. The same items are priced each month to the extent possible, and new items are substituted as the current items disappear from the market.¹ The outlet sample is rotated approximately every four years, and new item samples are initiated at that time.

It has long been thought that scanner data can be of benefit to the CPI for some or all of the following reasons:

1. Scanner indexes promise greater precision or lower variance than traditionally collected price data.
2. Scanner data record transactions that have actually taken place, whereas the CPI collects the prices of items on the shelf whether transactions actually took place at these prices in the given month or not.
3. Scanner data result in indexes with finer “granularity”—greater commodity detail.
4. The scanner sample is more representative of the universe because the weights are estimated more accurately and updated more frequently.
5. The scanner data are cleaned according to rules that can be applied consistently and studied academically.
6. Scanner data present the opportunity to implement superlative indexes.

This paper explores one way that the CPI could use scanner data in place of traditionally collected price data.

ACNielsen and Information Resources are the only major U.S. vendors of supermarket scanner data, routinely collecting scanner data from the retailers (e.g., Kroger, Safeway, and Giant) and marketing them to the manufacturers (e.g., General Mills, Nabisco, and Pepsi). Although the manufacturers could collect these data themselves, there are clear economies of scale enjoyed by the vendors, and in fact no manufacturer collects these data directly. The vendors add value by sorting the data into useful categories (e.g., cereal), subjecting them to range and consistency checks, and making them available in a standard format. The data contain three principal dimensions—product, time, and geography—and any desired subset is available. Thus, BLS can theoretically request data for the items that are currently

1. Occasionally new outlets are chosen to augment the sample in cases in which serious attrition has occurred due to the closing of outlets.

priced (i.e., according to BLS item definitions) and also for the CPI index areas. This paper is a progress report on a CPI program initiative to construct scanner-based test indexes for breakfast cereal in the New York metropolitan area.

The CPI is calculated in two stages based on the BLS partition of the consumption universe into 211 item strata and 38 geographical index areas. In the first stage a basic index is calculated for each of the $211 \times 38 = 8,018$ index area-stratum combinations. The basic indexes for most item strata are constructed using a weighted geometric mean (Geomean) formula; the few remaining item strata use a modified Laspeyres formula. The second stage of the calculation uses the Laspeyres formula to combine appropriate sets of basic indexes to yield various higher level aggregates. These aggregates include the national all-items CPI along with intermediate-level aggregations, for example, national food and beverages, or geographical areas such as New York all-items. The CPI for the New York consolidated metropolitan statistical area (referred to in the CPI as A101) is the aggregation of three basic index areas, A109 (New York city), A110 (New York-Connecticut suburbs), and A111 (New Jersey-Pennsylvania suburbs).

To evaluate the possibility of improving the accuracy of the CPI by using scanner data, BLS created the ScanData team. ScanData's objective was to determine the feasibility of incorporating scanner data into the monthly CPI production process. The method has been to produce demonstration or test indexes for breakfast cereal in A101 New York using scanner data. A success in producing such indexes would confirm that it is indeed possible to improve the CPI using scanner data, whereas a failure would be evidence to the contrary. The goal is indexes that

1. Are produced on the CPI production schedule;
2. Cover the entire domain of a basic item or area stratum by combining or "amalgamating" scanner and CPI data to eliminate outlet and geographic gaps;
3. Are consistent with CPI sampling principles;
4. Are based on a sample that is rotated and refreshed at least as often as under the current CPI procedure;
5. Use both the standard CPI geometric mean formula and a superlative index formula;
6. Use data cleaned at least to current BLS standards; and
7. Use prices with tax.

ScanData has constructed test indexes based on Nielsen breakfast cereal data for New York monthly in real time since March 1998, and it has similar indexes based on back data in a somewhat different format from September 1994 to January 1998. These data have been used to construct Geomeans, Laspeyres, Törnqvist, and Sato-Vartia indexes. The fixed-weight indexes (i.e., Geomeans and Laspeyres) use weights based on the

previous calendar year, whereas the weights for the Törnqvist and the Sato-Vartia are updated monthly.

The results show that, over the whole sample period, March 1998 through December 2000, the A101 scanner Geomean index was 105.2 and the CPI was 99.5 on a February 1998 = 100 base. On average, the scanner indexes have less variability than the current CPI. There are about 80,000 scanner quotes collected each month in New York and about 55 traditional CPI quotes, and hence there is a potential reduction in the standard error by a factor of about $\sqrt{80,000/55} = 38.1$. In fact, this is quite a bit too optimistic, mainly because the scanner data allocate the sample very inefficiently from the CPI point of view. Another indication of the reduction in the standard error is the mean absolute percent error (MAPE) about the mean. If the level were constant from month to month, the MAPE would be a measure of the spread of the distribution. As it is, the MAPE includes both the percentage error and the change in the level. Through December 2000, the CPI Geomean relatives have a MAPE of 2.63 percent compared with 1.23 percent for the scanner Geomean relatives, an indication that the scanner relatives are more precise but not nearly by a factor of 38.1.

Section 2.2 discusses the majority of the known technical issues in the implementation of the scanner indexes. The unit values issues are deferred to section 2.3. The scanner data are used to compute a variety of indexes, the formulas for which are presented in section 2.4. Section 2.5 presents a sample of the results of the real time experiment from March 1998 to December 2000 and also the earlier data referring to October 1994 through January 1998. Some conclusions are presented in section 2.6.

2.2 Technical Issues

The production of scanner indexes has involved dealing with a host of technical issues, some of which have been considered in the academic literature, others of which have not. The technical issues other than unit values and the index calculation are considered here. A summary of the status of these technical issues is found in table 2.1, followed by a definition and a more detailed discussion. A “Y” in the “Envisioned” column indicates that ScanData has a solution to the issue in mind, whereas an “N” is an indication simply that there is work to do. The solution to a given issue is said to be “Designed” if there is a mathematical solution on paper and agreement on the appropriate procedure. A solution is said to be “Implemented” if it exists at present in the ScanData computer program. There are “Results” if data have been produced using the implemented solution.

2.2.1 Quote Timing

The CPI collects data for the first eighteen business days (i.e., Mondays through Fridays, excluding holidays) in the month, except for November

Table 2.1 Status of Scanner Issues

Issue	Envisioned	Designed	Implemented	Results
Quote timing	Y	Y	Y	Y
Refreshing sample	Y	Y	Y	Y
Quote eligibility	Y	Y	Y	Y
Sales taxes	Y	N	N	N
Amalgamation	Y	Y	Y	Y
Migrating quotes	Y	N	N	N
Data cleaning	Y	Y	Y	Y
Imputation	Y	Y ^a	Y ^a	Y ^a
Variances	Y	Y	Y	Y
Geography	Y	Y	Y	Y
Average prices	Y	N	N	N
Supply disruptions	Y	Y	Y	Y
Aggregation	Y	Y	Y	Y

^aDecision being reconsidered.

and December, in which months it collects data for the first fifteen days. Each CPI quote is assigned to one of three pricing periods of five or six business days each, and the data from all three pricing periods are used to produce a monthly index. Nielsen scanner data, by contrast, are collected weekly (a week is Sunday to Saturday) and ScanData receives monthly shipments with either four or five weeks of data. The last week in each shipment is the third week of the calendar month, and the data are due ten business days after the end of the third week. The data have almost always arrived two or three days early. The first week of the month is defined as the first week with at least five days in the given month, irrespective of any holidays. Thus, the designated “first” week of the month occasionally includes the last day or two of the previous month, but never more than this. Since the prices generally refer to a whole week, they are also the prices for the actual first week of the month. ScanData is able to use data for the first three weeks of the month in the test indexes and can produce the test indexes a few days earlier than they would have been needed for the regular CPI computation. The indexes compare prices in the first three weeks of the current month to the first three weeks of the previous month.

The result of these collection rules for calendar year 2000 has been that the median CPI collection period covers the first through the twenty-fifth of each month, and the median scanner collection period covers the second through the twenty-second. Therefore:

1. The CPI and the scanner indexes cover very similar time periods within the month.
2. The CPI data are collected a little later in the month.
3. The scanner data explicitly cover Saturday, Sunday, and holidays.

2.2.2 Refreshing the Sample

As noted above, the CPI outlet and item samples are rotated (the old sample is dropped and a new one takes its place) every four years in current practice. Nielsen refreshes (i.e., adds units to) the outlet sample periodically to maintain sample size and to ensure that it continues to reflect the market in terms of the distribution of outlets by geography, organization,² format type, size, and age. Since the scanner data consist of a census of the items in the category in the sampled stores, the scanner item samples are refreshed continuously. For the test indexes that use the Geomeans or the Laspeyres formulas, ScanData refreshes both the weights and the item sample each year using the expenditure patterns of the previous year. Thus the 1999 scanner weights for New York cereal are based on 1998 expenditures, the weights for 2000 are based on 1999, and so on.³

2.2.3 Quote Eligibility

Quotes in the CPI are eligible for pricing if they were selected in the most recent PPS sample. Scanner quotes are eligible simply if the items have been sold in the previous year. Once a quote is selected, the CPI collects its price if it is present on the shelf regardless of whether it has been sold at that price recently or at all. Scanner prices, on the other hand, are “transaction” prices: that is, there is a price if and only if the item was sold during the given week, regardless of whether it is on the shelf at any particular time. Consumer Price Index quotes are imputed if the item is not available for purchase when the CPI data collector appears. Scanner quotes are imputed if they are eligible and if there were no sales in the first three weeks of the month.

2.2.4 Sales Taxes

For cereal as for most other items, the CPI collects prices net of any applicable tax. The sales tax is applied subsequently, using a secondary source with the sales tax rates for all of the jurisdictions in which there are CPI outlets. Scanner prices are also collected untaxed. For confidentiality, however, the vendors will not disclose the exact location of the outlets, so ScanData cannot add the sales taxes in the same manner as for CPI data. This is not a problem here, because breakfast cereal is not taxed in the New York

2. The term “organization” is used rather than “chain” to denote what is normally referred to as a supermarket chain in order to avoid confusion with the technical term “index chaining.” In addition, some organizations that appear to be supermarket chains to the consumer actually consist of independently owned stores with a common logo.

3. One exception should be noted: the Nielsen data for the three New York index areas became available in the current format only in February 1998, and hence both the 1998 and 1999 weights for New York were based on eleven months of data rather than the desired entire year. In addition, the 1998 indexes were based on the 1998 weights rather than, as would have been preferred, the 1997 weights.

CMSA. ScanData has not designed a solution to this problem as yet. It is hoped that Nielsen can help with the taxes, with the caveat that Nielsen may not realize what a complex problem it is. A second best solution, which BLS can implement independently, is to calculate a population-weighted average sales tax each month for each item stratum in each index area and apply the average sales tax to all of the outlets in the index area.

2.2.5 Amalgamation

The current CPI is designed so that every item in every outlet in all of the urban areas of the United States has a chance of selection in accordance with the PPS methodology. The scanner database is a proper subset of the universe covered by the current CPI. Therefore the scanner index will not cover the appropriate universe unless the outlets not covered by scanner data (for example, because of their type—as in the case of independent grocers—or location) are covered by retaining existing CPI quotes.⁴

The indexes reported in this paper combine preliminary scanner indexes (using scanner data only) with indexes constructed from CPI data for the nonscanner universe, a procedure called amalgamation. In particular, from 1998 through 2000 in New York there were four CPI quotes not in the scanner universe, all of which were from mass merchandisers or wholesale clubs. The mass merchandisers will be covered in future Nielsen data deliveries. The results show that the amalgamation in the New York cereal indexes did make a difference, and given that the CPI must be unbiased, it is necessary.

The amalgamation uses CPI weights, which come from the POPS. For example, if the current CPI had ten quotes and nine of them were from outlets covered by scanner data, the price relative computed from the scanner data would receive the sum of the weights of the nine quotes, and the non-scanner quote would retain its original weight. To date the Geomean formula has been used for amalgamation. Since the Törnqvist and the Sato-Vartia indexes rely on current period quantities, these formulas cannot be used for amalgamation. In a hypothetical production mode, Laspeyres amalgamation would be used in the Laspeyres strata.

2.2.6 Migrating Quotes

The CPI objective is to track changes in the prices paid by the residents of a particular index area for items in a particular item stratum. The CPI outlet surveys reveal that consumers make some purchases in outlets located outside the consumers' home areas. Hence, some of the CPI quotes are not collected in the area to which the index refers but represent purchases of consumers from the given area in other places because of travel,

4. Since the scanner markets are based on television markets, they do not always contain the corresponding census index areas as proper subsets. On the other hand, Nielsen has done a good job of reconfiguring its data to the CPI areas.

mail order, or other factors. Scanner data refer to the prices charged by the outlets in an index area, which is not quite the same as the prices paid by the residents of the index area. Thus, a scanner index based on outlets in a given area will not represent the CPI objective. None of the current New York CPI cereal quotes has “migrated” out of the New York metropolitan area, although some of the current CPI quotes for index area A109 are collected in A110 and A111. For this reason, a more adequate treatment would involve making the index for A109 a weighted average of the indexes for all three index areas.

2.2.7 Data Cleaning

Before indexes can be computed, the data quality must be assessed and questionable observations deleted. This is an automatic process that corresponds to the current CPI procedure wherein a BLS economist (called a commodity analyst) investigates quotes with large price changes and makes decisions whether or not to use them in the index. Although Nielsen does implement quality checks, for purposes of the CPI they must be supplemented by BLS efforts. Therefore ScanData has developed a procedure for automatically cleaning the cereal data based on the following five rules:

1. Accept all quotes that do not decline more than 37.5 percent or increase more than 60 percent in a given month.
2. On quotes with a promotion on the lower price, accept all quotes that do not decline more than 60 percent or increase more than 150 percent.
3. Accept all quotes for which the elasticity implied by the two months is at least as large as the 1.0 implied by the Geomeans calculation. In other words, if the price goes down, the volume must increase more than proportionately (so that revenue does not decline), and conversely.
4. Accept all quotes if, in the last twelve months, the price has been as high as or higher than the current price, and if the price has been as low as or lower than the current price.
5. Do not use any quote that does not satisfy at least one of the preceding rules.

If the item is sold in multiple units, for example, as a “two-fer” or a “three-fer,” in one but not both of the relevant time periods, rules 1 and 2 are modified as follows:

- 1'. Accept all quotes that do not decline more than 50 percent or increase more than 100 percent in a given month.
- 2'. On quotes with a promotion on the lower price, accept all quotes that do not decline more than 68 percent or increase more than 212.5 percent.

Rule 5 applies in both cases: that is, a quote must pass only one of rules (1 or 1'), (2 or 2'), 3, or 4. It has been found that the application of these rules deletes a few questionable quotes but accepts most of the data.

2.2.8 Imputation of Missing and Suspect Prices

In both the current CPI and the scanner index, the cleaned database contains a record for each quote. The record includes (if available) the quote's collected price, its effective (per ounce) price, and its derived price. The derived price (also per ounce) is either the effective price, where acceptable, or the imputed price otherwise. As in the CPI, only quotes with both an acceptable price in the current month and a derived price in the previous month are used in scanner index calculation.

Missing data for the scanner indexes (as for CPI data) are imputed for use in future months, and unacceptable prices are treated as missing data. In this process, missing prices are imputed by moving the last acceptable price forward by the (chain of) stratum relative(s). This is equivalent to current CPI practice, according to which the last acceptable price is moved forward each month by the index area–stratum relative. Once the indexes have been calculated, the imputed prices are calculated explicitly and entered as derived prices in the database for use in future months.

In the CPI, if an item is not on the shelf, is not expected to be restocked shortly, and has not been sold in the last few days, it is reported as missing. In cases in which the item is not expected to return, the CPI field representative will select a replacement. In the scanner procedure, if an item were not sold in any of the first three weeks of the month, it would be treated as missing. In scanner data indexes, however, replacements for missing items are not sought, since the data already include all possible items.

There is a new issue that arises in full force with scanner data. Because of the 100 percent sampling rate for items within an outlet, new items appear in scanner data much more often than in the CPI; moreover, there is virtually no lag between the time a new item appears in the outlet and the time it is available for our use. Often, at least for cereal, the prices of goods newly introduced into the market are not equilibrium prices but rather test prices set by the cereal manufacturer's product manager in the hope of obtaining a relatively stable volume at an acceptable price. They are erratic, sometimes beginning quite low and then increasing to a level comparable to similar products, and at other times beginning at a comparable level, declining in a deep sale, and then returning to a comparable level. This process can be repeated for a single item over its first few months. It can be fascinating to observe the resulting price trajectories, but any information on inflation is swamped by the wide variations in the prices resulting from the marketing process. The CPI has long faced the similar problem of products that disappear after a dramatic price reduction or a "closeout special." Since in formulas with constant weights (such as the Geomeans) the long-run effect on an index of a product that appears and then, after some time, disappears is just the ratio of the first and last prices, the introduction and withdrawal of items from the market can have a nontrivial effect on the index. The solu-

tion to this problem is in an imputation system that is based on something other than the last acceptable collected price. Thus, in table 2.1 the imputation design is labeled as a decision being reconsidered.

2.2.9 Variances and Replicates

The CPI variances are computed on the basis of replicates.⁵ In the current CPI, the small sample sizes mean that there are just two replicates for most index areas to serve as the basis of an estimate of the sampling variance. Leaver and Larson (2001) used a stratified jackknife calculation based on a segmentation of the scanner sample into separate index area–organization-identified strata: one for each of three to eight major organizations and one for the remaining scanner outlets within each index area. Within each stratum they created clusters of from one to three outlets each, and a replicate consisted of the whole sample with one of these clusters deleted. In this way the much larger scanner sample could support 126 clusters with one replicate for each of them.

The result was a much more precise estimate of the variance. It turns out that the actual reduction in the standard error is by a factor of 7.0 for the one-month comparison and 5.7 over twelve months. This is quite a bit less than the 38.1 factor calculated above on the basis simply of the relative number of quotes. The reason for this divergence is that the CPI sample is optimized for the calculation of the CPI, whereas the scanner sample is optimized for quite different purposes. In particular, the scanner data samples in the larger organizations are relatively much larger than in the small chains and independents, and this small sample makes a great deal of difference in the variance estimates. Leaver and Larson discovered the importance of this factor. It is as if we reduced the variance of a part of the sample with half the weight to zero while leaving the rest alone. The variance of the whole sample would simply be reduced by half, whereas the standard error would be reduced by only $\sqrt{1/2}$, regardless of how much we oversampled.

2.2.10 Northeast and National Geography

The population target of the CPI is the noninstitutional population living in metropolitan or urban nonmetropolitan areas. As noted above, the CPI has partitioned the urban United States into thirty-eight index areas.⁶ Scanner data are available for “markets,” which are generally smaller than the U.S. Census–defined metropolitan areas. To reconcile the geography, Nielsen has mapped its entire U.S. sample into the CPI index areas. The result is coverage of all of the thirty-one self-representing metropolitan areas

5. See Swanson (1999). A replicate is a sample consisting of only part of the universe of available quotes. An estimate of the sampling variance can be obtained by comparing the indexes estimated from the replicates.

6. The index areas are thirty-one self-representing metropolitan areas and seven “region-size classes” for the remainder of the covered population; for details, see Williams (1996).

except Anchorage, where Nielsen does not collect data. All seven non-self-representing areas were covered, and all of the data for the smaller places were used. Thus, the coverage of the smaller places was more complete than in the CPI, since these places are currently sampled.

2.2.11 Average Prices

The CPI computes average prices for a number of items for the convenience of users. It is clear that extremely accurate average prices could be computed using scanner data. To date there has been no decision as to which average prices to compute—that is, whether to simply continue the current set or expand it.

2.2.12 Protection against Data Supply Disruptions

The CPI program collects most of the prices it uses; consequently, the data supply is controlled within the program, and the program is responsible for any data supply disruptions. Delegating the basic data collection to a scanner data vendor, however, creates the possibility of data supply disruptions that are not under the direct control of the CPI. Therefore, in order for the CPI to fulfill its responsibility for a continuous flow of data, there must be a backup. One backup could be to continue to process the POPS as it is currently done to obtain (a) the weights required for the scanner amalgamation procedure and (b) a sample of quotes that could be initiated as a fallback in case of need, including a vendor supply disruption. No supply disruptions have occurred while we have been purchasing scanner data on a flow basis, and none is expected. The effect of a backup would be to preserve the ability of BLS to begin collecting prices for the scanner items in case of need.

2.2.13 Aggregation

The three New York areas are being combined to the A101, New York Metro, level, using the current CPI weights and the current Laspeyres calculation.

2.3 Unit Value Issues

This section considers three technical issues: weekly unit values, item definition, and organization-level, as opposed to outlet-level, indexes. The reason for considering these issues together in a separate section is that they all touch on the unit value issue in one way or another. A unit value is a quantity-weighted average price of an item. One way to compute a unit value is to divide the revenue for the item by the number of units sold. Unit values are currently used for the basic area-item indexes in most countries; they are not used in the United States because a weight is assigned to each individual quote.

Unit values are frequently encountered in scanner data since, for example, if an outlet has two different prices on an item in a week, the reported price is often the revenue for the item divided by the number of units sold. On the other hand, it may be appropriate to use unit values to combine observations over several weeks in the same month, or over outlets, or over similar items. Since a unit value is not a price in that often no one pays the unit value exactly, there is some controversy over whether our calculations should use them as if they were prices. ScanData is using unit values to combine the weeks of the month, UPC codes with minor differences, and the outlets within an organization in a given CPI area.

2.3.1 Weekly Unit Values

ScanData combines data for the three weeks of the month using unit values: the quantity-weighted average of prices for a given item.⁷ The price relative of an item is its unit value (averaged over the three weeks) for the current month divided by its unit value for the previous month. ScanData then computes the Laspeyres, Geomean, and other price indexes using these price relatives of item unit values with the appropriate PPS weights. One condition that must be satisfied for this to make sense is that the quantities must all be measured in the same units, a condition that is clearly satisfied here.

Instead of combining these prices and quantities in a unit value, one could combine them using the chosen index formula, in the same way that prices and quantities of different items are combined. However, the unit value approach more accurately reflects the preferences of the shopper who searches out the lowest prices each week, and also the consumer who stockpiles during a particularly good special, but then purchases nothing until the next special (see Triplett 1999). Consider the problems that arise by *not* using unit values and considering purchases in different weeks as different goods:

1. The weeks are arbitrarily defined, starting as early as the penultimate day of the preceding month, and as late as the fifth day of the current month. Thus, a purchase made every month on a particular date (e.g., any day between the sixth and the eleventh) will sometimes be allocated to the first week of the month and sometimes to the second.
2. The commodities purchased in the different weeks seem to satisfy the same needs and desires on the part of the consumer.
3. Unit values are required at some level in order to construct an index at all.
4. Consumers who stockpile are not indexed correctly without unit values over the weeks.

7. An item is commonly but not always a given UPC in a given store; see the subsection on Item Definition below.

5. *Not* using unit values implies an inherent rigidity in consumer behavior since it is assumed that the items in each of the three weeks are unrelated and that the elasticity of substitution among them is zero.

6. Defining items with a finer granularity, as is the case if quotes in different weeks are treated as separate items, results in more missing data and more imputations.

Thus, real inaccuracies can be introduced by *not* grouping identical commodities using unit values, and there is a powerful argument for considering purchases in the different weeks of the month as the same good.

2.3.2 Item Definition

Occasionally a manufacturer will keep the product constant but will create a new package with a new UPC code, a process called “churning.” Alternatively, and more interestingly, new UPCs sometimes appear that involve only small changes in the package size or flavor (e.g., blueberry vs. raspberry) that almost always sell at the same price. Changes in the package size can be used to indirectly raise or lower the effective (per ounce) price while keeping the shelf price constant. Both to reduce attrition and to capture these indirect price changes, it is important that these small to non-existent variations in the product be grouped together into a single item for index calculation.

Each month Nielsen supplies a list of new UPC codes together with the respective sizes and product descriptions. These are compared to the UPC codes already in use, and, if in the analyst’s judgment the differences are sufficiently small that the products are interchangeable, they are combined into a single item. The quantity (in ounces) of these combined items is the sum of the quantities sold of the constituent UPCs, the expenditure is the sum of the expenditures, and the price is the average price or unit value per ounce.

2.3.3 Organization-Level Indexes

The CPI is based on the price of a given item at a given outlet at a given time. Scanner data are available at the outlet level, and outlet data can easily be aggregated up to the organization level, in which case the quote is the unit value of a given item in a given organization at a given time. The organization-level index is based on the unit value or average price of an item at the outlets in an organization.⁸ It may be that unit values across outlets are appropriate: “if individual outlet data on transactions were not avail-

8. Each month the scanner data come with codes for each organization, without identifying the particular organization. There is always an “Other” organization representing smaller stores and independents. Heretofore the organization level index has been computed treating the other stores as if they constituted an actual organization. It has been proposed to modify this by treating each of the other stores as if it was a separate organization.

able or were considered to be too detailed, then unit values for a homogeneous commodity over all outlets in a market area might form the lowest level of aggregation” (Diewert 1995, 22). In fact, “everyday life suggests—and the suggestion is confirmed by our coffee data set—that consumers easily switch between outlets in response to relative price changes” (de Haan 1999, 64). Thus it has been thought that aggregation across outlets in general is not particularly controversial, and aggregation across different outlets in the same organization would seem to be even less so. The extra detail provided by outlet-level data is not much trouble to process, but it turns out that organization-level data are less expensive than outlet-level data because there are only a fraction as many data points. Perhaps of more significance is the need to reflect adequately the shoppers who search out the best sale, on the one hand, as opposed to the ones who always shop at the same outlet.

So far, the organization-level and outlet-level indexes for cereal have been quite close. In a hypothetical production mode, budget constraints may make organization-level indexes appealing. Nevertheless, there are three compelling reasons to continue to receive outlet-level data in a research project:

1. Outlet-level data allow the study of the differences between organization-level and outlet-level indexes.
2. Outlet-level data allow ScanData more control of the data quality.
3. Outlet-level data facilitate variance estimation.

2.4 Index Formulas

The scanner project originally was intended to produce Geomean indexes comparable to the current CPI. However, the scanner sample leads naturally to superlative indexes since the quantity of each item is collected each month along with the price. Therefore, Törnqvist and Sato-Vartia formulas have also been computed, along with the Laspeyres. The status of the various alternative index calculations is summarized in table 2.2.

As noted above, the second stage of CPI calculation is aggregation of basic indexes. The CPI for A101 New York cereal is the aggregate of the three

Table 2.2 Status of Index Formulas

Index	Envisioned	Designed	Implemented	Results
Geomean	Y	Y	Y	Y
Monthly Törnqvist	Y	Y	Y	Y
Annual Törnqvist	Y	Y	N	N
Laspeyres	Y	Y	Y	Y
Sato-Vartia	Y	Y	Y	Y

basic indexes for A019, A110, and A111. With respect to any index formula X , the A101 index at t compared with base month b is calculated from what are referred to in CPI terminology as “cost weights.” The month t cost weight for A101 is aggregated from the three constituent index area cost weights,

$$(1) \quad C_X^t = \sum_{m=1}^3 C_{Xm}^t.$$

The cost weight for index area m at time t is the product of A_m , the population-expenditure weight for cereal in index area m , and all of the period-to-period index relatives since the base time b ,

$$(2) \quad C_{Xm}^t = A_m \prod_{p=b+1}^t R_{Xm}^{p,p-1}.$$

The index of the price change from $t - k$ to t with respect to index formula X is just the ratio of the relevant cost weights from equation (1),

$$(3) \quad X^{t,t-k} = \frac{C_X^t}{C_X^{t-k}}.$$

The calculation equation (1) minus equation (3) corresponds exactly to the current CPI.

2.4.1 Geomean

The calculation of the Geomean relative used in the current CPI is

$$R_{Gm}^{t,t-1} = \prod_j I_j^t \left(\frac{p_j^t}{p_j^{t-1}} \right)^{S_j^0}.$$

Here j indexes the quotes in index area m , I_j^t is the indicator function for the presence of quote j in month t , and the quotes are weighted by S_j^0 , the expenditure share of quote j at time 0, $S_j^0 = E_j^0 / \sum_i E_i^0$. If there is no price in month $t - 1$, an imputed price is used if it is available.

The scanner Geomean relative in index area m from month $t - 1$ to t is

$$R_{Gm}^{t,t-1} = \prod_{c=1}^{n_m} \prod_j I_j^t \left(\frac{p_j^t}{p_j^{t-1}} \right)^{S_{cj}^0}.$$

In equation (4) each index area m is composed of n_m organizations c , j refers to the different quotes within an organization, and the quotes are weighted by the expenditure share of quote j in organization c at time 0, except for 1998 the previous calendar year,

$$(5) \quad S_{cj}^0 = \frac{E_{cj}^0}{\sum_i E_{ci}^0}.$$

The price p_j^t is the unit value of commodity j in organization c at time t ,

$$(6) \quad p_j^t = \frac{\sum_{i=1}^{n_{mc}} P_{ij}^t Q_{ij}^t}{q_j^t},$$

where P_{ij}^t and Q_{ij}^t are the prices and quantities, respectively, of j at the different outlets at time t , $q_j^t = \sum_{i=1}^{n_{mc}} Q_{ij}^t$ and there are n_{mc} quotes in each organization c . The A101 cost weight, using equation (4) in equation (2) with $X = G$, and then substituting into equations (1) and (3), results in the Geomean index $G^{i,t-k} = 100C_G^t/C_G^{t-k}$.

2.4.2 Monthly Törnqvist

This is the chained Törnqvist calculation, the geometric mean of the item relatives with weights equal to the average expenditure shares of the current and preceding months. Thus, with $W_{ej} = (S_{ej}^{t-1} + S_{ej}^t)/2$,

$$(7) \quad R_{im}^{t,t-1} = \prod_{c=1}^{n_m} \prod_j I_j^t \left(\frac{p_j^t}{p_j^{t-1}} \right)^{W_{ej}}.$$

The A101 cost weight, using equation (7) in (1) – (3) with $X = T$ results in the Törnqvist index $T^{t,t-k} = 100C_T^t/C_T^{t-k}$. The monthly Törnqvist relative has been calculated and amalgamated with CPI data from March 1998 forward. This amalgamation perforce uses the Geomean formula since there are no monthly expenditure weights for the CPI quotes.

2.4.3 Annual Törnqvist

In this form of the calculation, the base for the Törnqvist would be one fixed month that is used for a whole year. The current and base periods are single months. Assume, for example, that January were chosen as the base month. In this case, the February relative would be calculated as in the monthly Törnqvist above. The March relative, however, would be the Törnqvist of the March to January prices as if the February data did not exist: that is, the time zero prices and expenditure shares would refer to January. This would give up the pretense of the discrete time approximation to the Divisia but would eliminate any monthly chaining bias.

2.4.4 Laspeyres Index

This is the textbook Laspeyres, without any correction for formula bias, using weights updated once per year just as the Geomean. Thus the Laspeyres relative is

$$(8) \quad R_m^{t,t-1} = \frac{\sum_{c=1}^{n_m} \sum_j I_j^t p_j^t q_j^0}{\sum_{c=1}^{n_m} \sum_j p_j^{t-1} q_j^0} = \sum_{c=1}^{n_m} \sum_j I_j^t S_{ej} \frac{p_j^t}{p_j^{t-1}},$$

where the shares S_{ej} are defined in equation (5). The A101 cost weight, using equation (8) in (1) – (3) with $X = L$, results in the Laspeyres index $L^{t,t-k} =$

$100C_L^t/C_L^{t-k}$. The Laspeyres provides an upper bound with which to test the other indexes.

2.4.5 Sato-Vartia Index

The Sato-Vartia relative is a logarithmically weighted geometric average of the price relatives in which the weights W_{cj} are proportional to $m_{cj} = S_{cj}^t - S_{cj}^{t-1}/\ln S_{cj}^t - \ln S_{cj}^{t-1}$, so that

$$(9) \quad R_{Sm}^{t,t-1} = \prod_{c=1}^{n_m} \prod_j I_j^t \left(\frac{p_j^t}{p_j^{t-1}} \right)^{W_{cj}}.$$

The A101 cost weight, using equation (9) in (1) – (3) with $X=S$, results in the Sato-Vartia index $S^{t,t-k} = 100C_S^t/C_S^{t-k}$. The Sato-Vartia amalgamation uses the Geomean formula since there are no expenditure weights for the CPI quotes. Along with the Fisher, the Sato-Vartia satisfies more of the statistical axioms describing a desirable index than any other does. However, like the Törnqvist, it does not satisfy the monotonicity axiom (Reinsdorf and Dorfman 1999).

2.5 Results

Scanner indexes were calculated for the three New York index areas on the basis of February 1998=100 through December 2000. These were aggregated—using the appropriate weights—across the three areas. Table 2.3 below gives the various scanner indexes' results along with the CPI for cereal for A101.⁹ As we see from table 2.3, by the end of this period, the A101 CPI had fallen to 99.5, a decrease of 0.5 percent over thirty-four months. Over the same period all of the scanner indexes, whether amalgamated or not, increased, and ended in the 104.9–107.8 range. In addition, there were several large month-to-month changes in the A101 CPI, especially in October 1999, for which there was no corresponding change in the scanner indexes; there is no obvious explanation save the CPI's small sample size. Looking ahead, we see that this pattern was much attenuated in the national cereal CPI.

Unless there is some systematic divergence in price change between the scanner universe and the CPI universe, the scanner indexes will give a more accurate measure of inflation. By the end of the thirty-four-month sample period, the A101 CPI was 5.5 percent lower than the scanner Geomean index and 5.7 percent lower than the amalgamated Geomean. The Törnqvist, whether amalgamated or not, showed a bit less inflation than the Geomean,

9. The BLS does not publish an index for breakfast cereal below the national level. The New York cereal indexes reported here are unpublished, primarily because of the small sample sizes on which they are based and the resulting high variances. The comparison is somewhat muddied because the CPI switched from the Laspeyres to the Geomean for the lower level cereal indexes in January 1999.

Table 2.3 Organization-Level Scanner Indexes, 1998–2000

Month	A101 CPI	Geomean		Törnqvist		Sato-Vartia	
		Scanner	Amalgamated	Scanner	Amalgamated	Scanner	Amalgamated
		March 1998	98.8	100.7	100.7	100.7	100.8
April 1998	99.6	100.2	100.2	100.2	100.4	100.4	100.4
May 1998	99.9	100.5	100.5	101.1	101.0	101.0	101.0
June 1998	101.6	100.2	100.1	98.4	99.2	99.2	99.2
July 1998	102.1	99.8	99.8	99.1	99.5	99.5	99.5
August 1998	101.4	100.2	100.2	100.1	100.4	100.4	100.4
September 1998	99.6	97.6	97.7	97.0	97.5	97.5	97.6
October 1998	101.1	98.3	98.4	97.0	97.8	97.8	97.8
November 1998	101.4	100.4	100.4	100.8	101.3	101.3	101.2
December 1998	102.1	100.8	100.8	101.1	101.8	101.7	101.7
January 1999	103.3	100.6	100.7	100.1	100.5	100.7	100.7
February 1999	102.5	100.0	100.1	97.0	98.4	98.5	98.5
March 1999	104.9	100.2	100.4	97.0	98.6	98.6	98.8
April 1999	102.1	100.7	100.8	98.5	100.1	100.1	100.3
May 1999	101.7	102.1	102.2	100.5	101.8	101.9	101.9
June 1999	104.7	101.8	102.2	98.3	100.4	100.9	100.9
July 1999	103.8	103.0	103.5	100.3	102.2	102.7	102.7

August 1999	103.2	102.8	103.2	100.6	101.1	102.3	102.8
September 1999	104.6	100.3	100.9	98.4	99.1	100.0	100.6
October 1999	99.1	102.9	103.3	100.5	100.9	102.8	103.2
November 1999	98.7	105.3	105.6	105.7	105.9	107.3	107.5
December 1999	100.9	104.6	105.0	102.3	102.7	104.9	105.2
January 2000	97.5	103.0	103.5	101.3	101.9	103.0	103.6
February 2000	102.8	101.9	102.5	98.4	99.1	100.9	101.5
March 2000	100.7	102.9	103.3	100.4	100.9	102.7	103.1
April 2000	104.0	104.9	105.2	102.5	103.0	105.2	105.6
May 2000	104.7	105.2	105.5	102.9	103.3	105.4	105.7
June 2000	101.1	103.9	104.3	99.0	99.5	102.4	102.8
July 2000	100.0	104.1	104.4	100.8	101.2	103.6	103.9
August 2000	101.6	104.1	103.9	101.3	101.3	104.1	103.9
September 2000	98.6	100.6	101.2	98.2	98.9	100.7	101.3
October 2000	102.7	103.1	103.6	100.3	100.9	103.3	103.8
November 2000	99.7	106.7	107.0	106.4	106.9	109.2	109.4
December 2000	99.5	105.0	105.2	104.9	105.1	107.7	107.8

whereas the Sato-Vartia showed quite a bit more. The very slightly less Törnqvist inflation, compared with the corresponding Geomean, cannot plausibly be taken as evidence that the elasticity of substitution is less than the 1.0 assumed by the Geomean. On the other hand, the divergence between the Törnqvist and the Sato-Vartia may lead us to question the degree to which the latter is “nearly superlative.”

The amalgamated indexes from table 2.3 are plotted against the CPI in figure 2.1. The CPI used the Laspeyres formula, an upper bound to the true cost-of-living index, through December 1998, and the A101 CPI shows more inflation than any of the scanner indexes through September 1999 before dropping convincingly below.

The corresponding scanner indexes (i.e., those that are not amalgamated) are plotted in figure 2.2 below. There is no particular reason to expect that the Geomean, Törnqvist, and Sato-Vartia should bear any particular relationship to each other. It is interesting that once again the superlative Törnqvist and the nearly superlative Sato-Vartia do not track each other particularly closely. In fact, the Törnqvist was below the other two for seventeen of the twenty-two comparison months, including all of the last thirteen, before almost catching up at the end to finish slightly below the Geomean.

Figure 2.3 below considers the scanner indexes without amalgamation




Fig. 2.1 Organization-level amalgamated indexes and the CPI, 1998–2000




Fig. 2.2 Organization-level scanner indexes, 1998–2000




Fig. 2.3 Organization-level scanner indexes, 1994–2000

Table 2.4 **Relatives by Month**

Month	Product of Relatives		Ratio	Ratio per Month
	Geomean	Törnqvist		
January	0.918	0.886	1.036	1.006
February	0.983	0.932	1.055	1.009
March	1.051	1.049	1.002	1.000
April	1.065	1.106	0.964	0.994
May	0.978	0.972	1.006	1.001
June	0.932	0.861	1.083	1.013
July	0.990	1.019	0.972	0.995
August	1.020	1.044	0.977	0.996
September	0.848	0.845	1.003	1.001
October	1.141	1.155	0.988	0.998
November	1.087	1.166	0.932	0.990
December	1.044	1.049	0.995	0.999
Total	1.0219	1.0181	1.0037	1.0027

back to September 1994, the beginning of the cereal data available to Scan-Data.¹⁰ Here the final ranking of the formulas is the same as in the more recent period, and once again the Törnqvist is below the others during most of the later period, despite almost catching up at the end. In fact, although the Törnqvist was ahead of the Geomean for seventeen of the first twenty months, through May 1996, it is ahead for only two of the last thirty-one months, beginning in June 1998.

Some understanding of this unexpected temporal relationship between the Geomean and the Törnqvist can be found in an unanticipated seasonal pattern. The ratio of the Geomean to the Törnqvist in December 2000 from table 2.3 is 1.00372: the Geomean has outpaced the Törnqvist by about 0.37 percent over seventy-five months. However, the relationship between the indexes varies substantially over the months of the year. Table 2.4 presents the product of the Geomean and Törnqvist relatives by month over the whole sample period. The cell in the January row and the Geomean column is the product of the relatives for the six Januaries in the sample, and so forth through September. The entries for October through December are the product of the seven relatives for these months. The column labeled "Ratio" is just the Geomean entry divided by the corresponding Törnqvist entry.

10. Figure 2.3 is subject to the caveat that the format of the scanner data changed between January and February 1998 in a way that caused a break in the series. This problem was resolved here by arbitrarily making all of the scanner relatives for February compared with January 1998 equal to the corresponding national CPI relative. Note also that (except for the A101 CPI) the data in figure 2.3 are the scanner indexes *not* amalgamated with any CPI data. The reason for this is that the data from September 1994 through January 1998 were not received in real time, and hence it was not possible to amalgamate the CPI quotes not covered in the scanner universe.

“Ratio per Month” is either sixth or seventh root of the ratio—that is, the geometric mean of the corresponding ratio. The totals at the bottom are just the products of all of the entries in the given column.

It is clear from table 2.4 that there is a distinct seasonal pattern in the relatives, with the Geomean gaining during the first half of the year, January through June, and the Törnqvist catching up during the second half, especially during November. The reason that the Törnqvist tended to lead for the first part of the sample period is that the first comparison month was October, and hence the Törnqvist was ahead at the start due to this seasonal effect. Nevertheless, the Törnqvist tended to fall behind over time and ended, as mentioned, 0.37 percent behind the Geomean. The total in the “Ratio per Month” column shows that over an average year the Geomean gained 0.27 percent, and hence, had we corrected for the seasonal effects the Geomean would have ended $1.00272^{6.25} - 1 = 1.71$ percent ahead of the Törnqvist. A more direct calculation confirms this general idea and gives an even larger difference between the indexes. The annual relatives of both the Geomean and the Törnqvist indexes were calculated, and also the ratios between them. This eliminated eleven observations and reduced the sample to sixty-four annual comparisons. The geometric mean of the sixty-four annual relatives was 1.00554, a gain of 0.55 percent per year for the Geomean and of 2.99 percent over the whole sample.

These differences are quite striking since both the prices and the form of the index formula were the same, the only difference being in the weights. It would appear that this is a manifestation of the “high-frequency” phenomenon discussed by Feenstra and Shapiro (F&S) in this volume. In the F&S data the Törnqvist (their “chained Törnqvist”) resulted in much higher estimates of inflation than the Geomean, the opposite of the findings here. Feenstra and Shapiro found that often an item went on sale a week or two before ads appeared and that the ad appeared the last week of the sale. The result was that the increase in price at the end of the sale was weighted more heavily than the previous decline, and hence the Törnqvist increased too rapidly.

Feenstra and Shapiro present a sample data set of two items over six months in which the beginning and ending prices are the same for both items, and hence the ending index should be equal to 100. The Törnqvist index, however, stood at 154.1 at the end, whereas summing the data into months as in the ScanData methodology resulted in an ending Törnqvist of 95.2.¹¹ Thus, the monthly Törnqvist gave a low measure, whereas the weekly Törnqvist gave a high measure based on the same data. It would appear that item volume the month before a sale tends to be normal, but vol-

11. This is slightly different from the 145.2 figure given by F&S in their table 3, even after taking account of their initial index value of 92.7, due to rounding. The basic data are in their table 1.

ume after the sale tends to be suppressed due to stockpiling. Feenstra and Shapiro also present a “fixed-base Törnqvist” index, which ends at 103.0. However, since the beginning and ending prices were equal, none of these is correct. The evident inventory effect invalidates the assumption of static utility maximization underlying the Törnqvist formula. The F&S “fixed-base Törnqvist” formula is one attempt to deal with this, but clearly more work is required before such a formula can be used in the National Income Accounts.

In figure 2.4 below we plot the national CPI for breakfast cereal against the A101 CPI and the scanner Geomean indexes at the outlet and organization levels. The organization-level and outlet-level Geomean indexes track quite well, the organization level index trailing by only 0.38 percent after sixty-eight months, an indication that the unit value at the organization level results in only a small correction. The national CPI and the A101 CPI do not track each other very well at all, however, and hence, despite possible differences between the national and New York markets, we must suspect considerable small sample variability in the A101 CPI. In particu-




Fig. 2.4 Scanner Geomean indexes and the CPI, 1994–2000

lar the pronounced October 1999 dip in the A101 CPI is not reflected at all in the national CPI. By December 2000 the national CPI is above the New York scanner indexes, whereas the A101 CPI is far below them.

2.6 Conclusions

There have been numerous suggestions that scanner data could be used in the CPI for additional commodity detail. It has also been hoped that scanner data could result in more accurate low-level indexes and also provide the opportunity to implement superlative indexes at the lowest level. Another possible advantage is that the scanner sample may be more representative of the relevant universe because the weights are estimated more accurately and updated more frequently. Yet another advantage is that scanner data can be cleaned according to rules that can be applied consistently and studied academically. The purpose of the ScanData project was to study these propositions.

ACNielsen has been an excellent vendor, supplying the scanner data on time and dealing with issues as they arose. A computer program has been developed and extensively acceptance-tested to construct scanner-based test indexes for breakfast cereal in the New York metropolitan area. The technical issues have been largely dealt with, as we saw in sections 2.2–2.4. The data have been received on a schedule such that they could have been used in the CPI, the scanner sample has been refreshed appropriately, and the standards for quote eligibility are more closely related to the relevant theory than are those of the current CPI. The scanner data have been amalgamated with the few CPI quotes from nonscanner outlets to produce indexes that cover the CPI target universe. The automated computer data-cleaning system has been found to work adequately. Missing prices have been imputed using rules similar to those of the current CPI, and in the process defects have been discovered in those rules. The more accurate variance calculations that are possible with scanner data have been carried out. The scanner indexes have been aggregated to the appropriate higher levels and some protection has been found against supply disruptions in a hypothetical production environment. Issues with respect to average prices and taxes remain, but it would appear that these can be solved with a bit more effort. Issues revolving around unit values have been dealt with, and the index formulas have been implemented.

The results of Leaver and Larson (2001) show that the standard errors of the scanner indexes are about one-sixth of those of the current CPI, a significant reduction but not nearly what would have been expected simply on the basis of the increase in sample size. The calculation of the superlative Törnqvist index has revealed unforeseen problems with the superlative methodology. Although it has long been known that stockpiling and search behavior cause problems for price indexes, they are of such magnitude as to

preclude the traditional superlative indexes' being used with this fine a granularity in the product, time, and geographical dimensions. Although the use of more aggregated data would mitigate the problems of the superlatives, they are so fundamental that it is not clear that they would disappear at any level of aggregation.

References

- Boskin, Michael J., Ellen R. Dulberger, Robert J. Gordon, Zvi Griliches, and Dale Jorgenson. 1996. *Final report of the Commission to Study the Consumer Price Index*. U.S. Senate, Committee on Finance. Washington, D.C.: U.S. Government Printing Office.
- Bradley, Ralph, Bill Cook, Sylvia G. Leaver, and Brent R. Moulton. 1997. An overview of research on potential uses of scanner data in the U.S. CPI. In *Proceedings of the third meeting of the International Working Group on Price Indices*, ed. Bert M. Balk, 169–82. Voorburg, the Netherlands: Statistics Netherlands.
- Conference Board, The. 1999. *Measuring prices in a dynamic economy*.
- Diewert, W. Erwin. 1995. Axiomatic and economic approaches to elementary price indexes. Discussion Paper no. 95-01. Vancouver, Canada: University of British Columbia, Department of Economics.
- Haan, Jan de. 1999. *Empirical studies on Consumer Price Index construction*. Voorburg, the Netherlands: Statistics Netherlands.
- Haan, Jan de, and Eddy Opperdoes. 1997a. Estimation of the coffee price index using scanner data: The choice of the micro index. In *Proceedings of the third meeting of the International Working Group on Price Indices*, ed. Bert M. Balk, 191–202. Voorburg, the Netherlands: Statistics Netherlands.
- . 1997b. Estimation of the coffee price index using scanner data: The sampling of commodities. In *Proceedings of the third meeting of the International Working Group on Price Indices*, ed. Bert M. Balk. Voorburg, the Netherlands: Statistics Netherlands.
- Leaver, Sylvia G., and William E. Larson. 2001. Estimating variances for a scanner-based consumer price index. *Proceedings of the American Statistical Association, Government Statistics Section*.
- McKaig, S. Kate. 1999. Universal Product Codes. Washington, D.C.: CPI Tech. Memo, April.
- Reinsdorf, Marshall B. 1997. Constructing basic component indexes for the U.S. CPI from scanner data: A test using data on coffee. BLS Working Paper no. 277. Washington, D.C.: Bureau of Labor Statistics.
- Reinsdorf, Marshall B., and Alan H. Dorfman. 1999. The Sato-Vartia index and the monotonicity axiom. *Journal of Econometrics* 90:45–61.
- Swanson, David. 1999. Variance estimates for changes in the Consumer Price Index, January 1998–December 1998. CPI detailed report. Washington, D.C.: Bureau of Labor Statistics, December.
- Triplett, Jack E. 1999. Should the Cost-of-Living Index provide the conceptual framework for a consumer price index? *Economic Journal* 111 (472): F311–34.
- Williams, Janet. 1996. The redesign of the CPI geographic sample. *Monthly Labor Review* 119 (12): 10–17.

