Constructing High Frequency Price Indexes Using Scanner Data

Daniel Melser

School of Finance & Economics

University of Technology, Sydney

Haymarket, NSW 2007

Australia

Email: daniel.melser@uts.edu.au

Phone: (+61 2) 9514 4078

Draft: April 27, 2011*

Abstract: The advent of scanner data was greeted with great fanfare in the price index fraternity given the wealth of information it can potentially provide for use in index construction. Unfortunately, the integration of the scanner data into the construction of official indexes since this time has been somewhat underwhelming. One of the major factors behind the hesitancy of statistical agencies to use scanner data is the volatility in both prices and quantities of individual items. Much of this is to do with the semi-regular rotation of products on and off sale, and the very strong changes in purchase volumes that this induces. Inventory effects—which sees consumers restock during sales—can lead to asymmetric changes in quantities when compared with price movements. Standard index number methods can produce questionable answers in such circumstances. As a resolution to this problem we outline a stochastic mixture model of prices which we use in constructing price indexes. This model takes account of the step-like nature of prices, as products go on and off sale, and produces an index which we argue is both bias free and likely to have lower variance than alternative approaches.

Keywords: Price index, scanner data, sales, mixture model.

JEL Classification Codes: C14, C43, E31.

* This paper was prepared for discussion at the 12th Ottawa Group meeting, May 2011, Wellington, New Zealand.

1 Introduction

There was considerable excitement, and a flurry of research activity, following the 'discovery' of scanner (or barcode) data with regard to its uses in constructing price indexes (Silver, 1995, 1999; Silver and Webb, 2002; Silver and Heravi, 2001; Reinsdorf, 1999; Feenstra and Shapiro (eds.), 2003; Melser, 2006). In modern retailing almost all purchases of goods are scanned at the checkout and the barcode, the price and sales volumes are recorded. This information is then used by the firm in supply chain management and often passed on to a third party, such as a market research company, for analysis. Perhaps as much as one quarter or a third of all consumers' expenditures are recorded in this way.[1] On the face of it the potential of scanner data is enormous. Traditional statistical agency sampling methods involve visiting shops at specific times and obtaining isolated price quotes with no expenditure information. By contrast scanner data offers effectively a census of transaction prices along with associated quantity data. Given this it was hoped that scanner data would not only lead to more accurate measures of prices change but also more stable estimates given the wealth of data relative to the traditional sample.

However, now more than a decade has passed and unfortunately the extent of the integration of scanner data into the construction of official statistics has been disappointing. In the US the BLS makes no mention of its use of scanner data though has undertaken indepth research on the topic (see Richardson (2003) and the references included). The situation is similar in Canada, Australia, the UK and elsewhere. The only statistical agencies to make substantial use of scanner data in any meaningful way are Statistics Netherlands, Statistics Norway and the Swiss Federal Statistical Office—though Statistics New Zealand uses scanner data for weighting purposes. The Dutch have led the globe in their use of scanner data but it has not been trouble-free (Statistics Netherlands, 2003). Few have followed. This is despite exhortations by the Boskin Commission in the late-

[1] For example examining the breakdown of weights in BLS (2009) we can get a rough guide as to the proportion of expenditure which scanner data exists. For urban workers it is likely that the 7.746% of expenditure on *Food at Home* will be for scanned products; much of the 4.612% of spending on *Household Furnishings and Operation*, for items such as furniture, appliances, and tools, will be recorded electronically; as will the 3.663% of expenditure on *Apparel*; spending on gasoline of 4.429% is also likely to be collected electronically. If we include the spending on recreational goods, such as audio visual equipment and sporting goods as well as spending on cigarettes and tobacco and personal care goods this takes the count well above 25%. The proportion is likely higher than this and will no doubt expand in coming years.

1990s, with regard to the United States CPI, to explore the use of scanner data to to reduce the costs of data collection, help in estimating expenditure weights and to introduce new products in a more timely fashion (Boskin, Dulberger, Gordon, Griliches and Jorgenson, 1997, p.81). More recently the National Research Council (2002) report emphasized the potential advantages of the scanner data but sounded a somewhat more cautionary note about the care that was required in effectively using it in price indexes. The increasing circumspection with regard to scanner data reflects much of the research over the past decade which has had trouble constructing price indexes using such data. In an influential early study of Chicago coffee data Reinsdorf (1999) showed that under various 'reasonable' approaches he obtained estimates of price change varying between 71.5% and 371.1% over just two years.[2] Much of the other work on scanner data, such as by Silver and co-authors (Silver, 1995, 1999; Silver and Webb, 2002; Silver and Heravi, 2001), showed that different approaches could lead to staggeringly different answers. This is of course the opposite of what statistical agencies seek as they want to develop robust approaches which depend little on arbitrary assumptions. This has emphasized a point made by Triplett (2003) and others, that the standard approach to cost-of-living construction is insufficient when faced with the reality of scanner data.

In particular, we argue that there are two major issues which crop up in the construction of price indexes using scanner data. These are sales and stockpiling. At the level of individual prices there is extreme variability in prices at a daily, weekly, or even monthly frequency as a result, not of inflation, but of the sales and the discounting cycle. Consumers who take advantage of these sales can save significant amounts. Generally speaking the existence of two price modes for a good—the regular price and the sale price—poses problems for classical index methods which focus on the mean. We develop this point below and suggest alternative ways of modeling price change in this case. The existence of a sales cycle, and the fact that many goods are storable, leads to significant stockpiling activity during sales, with purchases skyrocketing, while during regular pricing demand is considerably weaker. The conjunction of significant fluctuations in price, stockpiling, and the separation of purchase and consumption poses significant problems for the construction of cost-of-living indexes. While discussion of sales behaviour has been a focus of attention in some quarters—usually in the marketing literature (see for example Van

[2]The strong price rise, even at the lower end of the spectrum, reflected the fact that there were significant increases in the price of coffee at this time as a result of supply-side factors.

Heerde, Leeflang and Wittink, 2000)—there has been little attempt to model this for the purposes of constructing a price index. Inventory behaviour by consumers has recently received significant attention (Hendel and Nevo, 2006a, 2006b) but the implications for price index construction are yet to be fully explored.

This paper address the broad question of how to construct unbiased and stable price indexes using scanner data at high frequency given the reality of sales and stockpiling. We seek to model the sales cycle and integrate this into a framework for constructing price indexes. While our approach is a little speculative, our objective, driven by the demands of official statistical agencies, is to seek a methodology which produces a unbiased estimate of price change with as little volatility as possible. In doing this we analyze a major US scanner data set made available for research purposes by Information Resources Incorporated (IRI) (see Bronnenberg, Kruger and Mela, 2008). The data involves weekly sales prices and quantities at the barcode level by store and metropolitan area from 2001 to 2006. The data is very rich and will be used extensively to illustrate our points.

In the next section we begin by addressing the problem of sales. We outline a model of price generation which enables the identification of sales and apply this to the IRI data set. Section 2 also examines the significance of stockpiling for our data. In section 3 we outline ways of thinking about the cost-of-living index in the context of this issue. Our approach suggests a number of practical ways of constructing price indexes on high frequency data which are likely to be bias-free and have relatively low variance. Section 4 compares the various approaches to price index construction empirically using the IRI data. We conclude by providing some summary comments in section 5.

2 Sales and Stockpiling

The use of sales discounts by retailers is ubiquitous. Either in large mega-retailing establishments or in smaller corner stores. While the classic reason for sales is to get rid of unwanted store inventory there are likely to be many other reasons too. Varian (1980) argued that sales enabled retailers to discriminate across informed and uninformed consumers, Sobel (1984) saw it as a way to occasionally sell to low reservation price consumers. There are a range of other reasons as well, such as the loss-leader strategy (Hess and Gerstner, 1987).

Despite the prevalence of sales, their modeling and identification has received relatively little attention. The standard approach being to define some level of price cut below which the good is identified as being on sale (Hendel and Nevo, 2006a; Feenstra and Shapiro, 2003). While likely to give a reasonable approximation the choice of cutoff appears somewhat contestable, and is likely to vary by product category and item. Feenstra and Shapiro (2003) chose a 5% cutoff while Hendel and Nevo (2006a) undertook their analysis for a range of cutoff values; 5%, 10%, 25% and 50%.

In Figure 1 we depict a price path for various types of laundry detergent. The outstanding feature of the behaviour of price is its oscillation between two different levels. That is, it is inherently bimodal in nature. Hence, we argue that in order to understand and measure price change it is likely to be necessary to have a model which incorporates the bimodal nature of prices. A natural way of modeling such a variable is through the use of a mixture distribution.

Figure 1: The Sales Cycle — Laundry Detergent Prices

(a) Product 9, Store 68

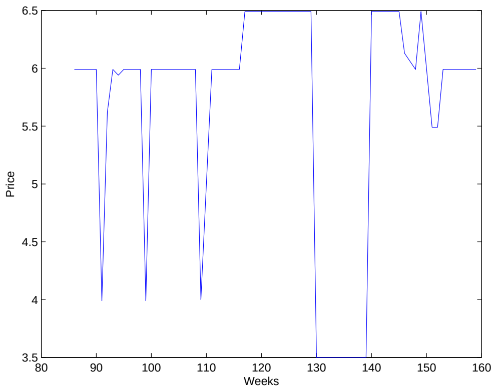

(b) Product 18, Store 1

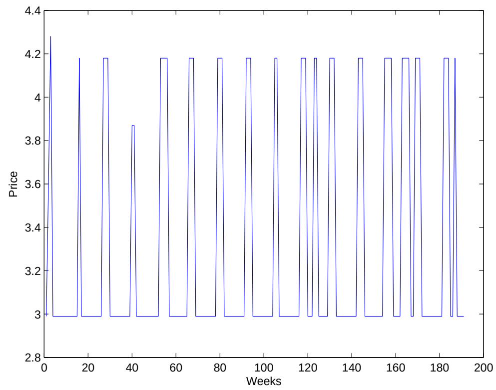

(c) Product 18, Store 21

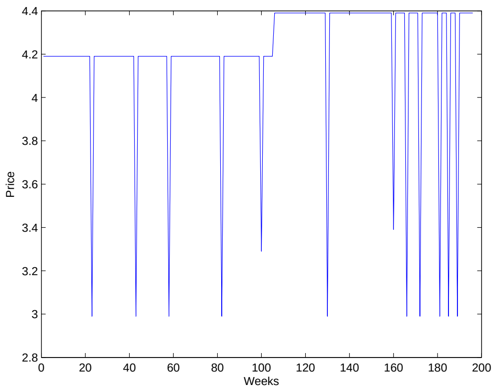

(d) Product 17, Store 59

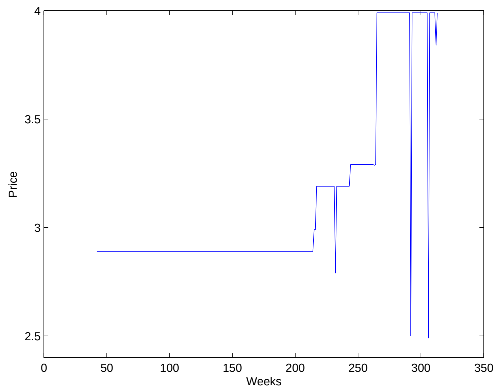

We consider a model of prices which explicitly accounts for the sales cycle. Let $\ln p_{ist}$ denote the log of price from store $s \in S_t$, where S_t is the set of stores open in time $t = 1, 2, \ldots, T$, for some good $i \in I_{st}$, where I_{st} denotes the set of products available in store s in time t. Now the price will will either be a draw from the 'regular' (higher mean) price distribution or it will come from the (lower mean) sales distribution. More formally we suggest the following stochastic mixture model for log prices,

$$\ln p_{ist} \sim z_{ist}\mathrm{N}\left(\alpha_{ist} - \beta_i, \sigma_i^2\right) + (1 - z_{ist})\mathrm{N}\left(\alpha_{ist}, \sigma_i^2\right), \ \beta_i \gg 0 \tag{1}$$

$$z_{ist} \sim \mathrm{B}(\omega_{it}) \tag{2}$$

Here α_{ist} can be thought of as the 'regular' price for good i in store s at time t. This is item, store and time specific. Potentially a number of different techniques could be used to estimate this, such as a polynomial functional form in time or a spline. Our approach was to model it using a nonparametric local regression method with a predetermined span. This performed relatively well in practice. However, note that a particular price, $\ln p_{ist}$, may also be a sale price. In this case the mean falls by the sale effect β_i where $\beta_i \gg 0$. This is incorporated into the sales price distribution. The variance of prices under these two regimes is assumed to be the same σ_i^2. The essential idea of this model is shown in Figure 2 below.

Figure 2: The Distribution of Prices

The key part of the model is that we also have the variables z_{ist}. These may be regarded in some sense as missing data and indicate whether a good is on sale $z_{ist} = 1$ or not on sale $z_{ist} = 0$. These are Bernoulli random variables distributed with probability

ω_{it}—the probability that a good is on sale. Note that this can change over time and reflects the propensity that stores have to discount. The parameters to estimate in this model are, excluding the subscripts, the various α, β, σ, ω and z variables. This model can be estimated iteratively using the EM (expectations maximization) algorithm (Dempster, Laird and Rubin, 1977). Here we effectively treat the sales indicators as missing data. A more detailed discussion of estimation is contained in the Appendix 7.1.

It should be noted that the model above is not a standard hedonic model. It is much more flexible than that. Indeed the output of the model is essentially a price index for each product in each store, α_{ist}. Clearly some additional approach is required to construct a price index across products. Our approach here is not (necessarily) to use a regression method but to focus in on identifying sales behaviour and use this in informing our price index methodology. In order to accurately identify sales we have intentionally kept the model very general.

In principle, our point is that in the presence of such a bimodal distribution the mean of the distribution is not going to necessarily be the best measure of it. In fact, it is straightforward to show—as we do in Appendix 7.2—that the variance of the mean is equal to,

$$\mathrm{Var}(\ln p_{ist}) = \sigma_i^2 + \omega_{it}(1 - \omega_{it})\beta_i^2 \tag{3}$$

Clearly, this is greater than σ_i^2, the variance of each of the individual mean parameters. The wedge between these two variances depends upon the frequency of sales and the size of the discount. Given this there is the possibility of reducing the variance of some estimator of the mean price change by explicitly taking account of the bimodal nature of prices.

In order to explore the feasibility of the mixture model we apply it to the IRI data set for a range of products; beer, carbonated beverages, coffee, deodorant, laundry detergent, milk, mustard and ketchup, paper towels, peanut butter, salty snacks, shampoo, soup, spaghetti sauce, sugar substitute, toilet tissue, toothpaste, yoghurt. The IRI data set is very large indeed and we focus on just a small component of it. We utilize data for just one of the available markets, Boston, over the whole 6 year period. In order to keep the size of the data to manageable proportions we have trimmed the selection of items to those that were available for at least a year and removed some at random stores. Some summary statistics on the data used are contained in Table 1 below along with estimates of the key parameters. In particular there is strong evidence of very significant sales behaviour. The average probability of a sale for the product categories varies significantly. If we calculate the mean percentage of sales across the data set this varies from just 8.46% for beer to 23.67% for spaghetti sauce. There is even more variability in the discount rate across products, i.e the β. For each item within a product category we calculate a separate discount. If for each product category we take the mean of these across all items then this yields an average discount which is lowest at 12.87% for beer and highest at 34.43% for soap.

Table 1: Summary Statistics for IRI Data — Boston

Product	Number of:			Expenditure ($)		Probability	Average Sales
	Items	Stores	Observations	Total	On Sale (%)	of Sale (%)	Discount (%)
Beer	558	33	175,658	10,773,410	2,831,948 (26.29)	8.46	12.87
Carbonated Beverages	863	9	355,071	14,401,327	5,606,342 (38.93)	21.79	25.51
Coffee	460	14	304,416	9,782,830	3,383,628 (34.59)	21.99	27.65
Deodorant	551	14	313,495	1,873,691	329,900 (17.61)	15.70	36.55
Diapers	260	37	248,830	8,624,339	2,130,034 (24.70)	20.04	24.32
Laundry Detergent	325	16	210,504	8,748,413	3,221,216 (36.82)	17.24	31.80
Milk	282	17	321,594	48,326,059	8,291,286 (17.16)	19.52	19.56
Mustard and Ketchup	225	23	245,953	6,051,455	1,288,858 (21.29)	17.71	25.97
Paper Towels	219	38	234,434	21,202,529	6,074,663 (28.65)	12.25	28.98
Peanut Butter	116	35	270,964	9,242,605	2,451,277 (26.52)	17.45	25.23
Salty Snacks	963	9	307,746	8,788,534	2,373,876 (27.01)	18.35	27.55
Shampoo	551	14	299,967	2,326,703	4,59,215 (19.74)	15.70	30.27
Soup	552	9	264,729	4,612,889	1,688,499 (36.60)	22.41	34.43
Spaghetti Sauce	422	12	288,320	5,465,383	1,998,879 (36.57)	23.67	28.22
Sugar Substitute	52	51	191,022	4,435,557	656,878 (14.81)	13.31	22.85
Toilet Tissue	163	32	241,574	22,487,510	6,322,188 (28.11)	13.42	27.97
Toothpaste	403	14	246,966	2,952,563	589,022 (19.95)	15.70	31.27
Yoghurt	550	6	272,707	10,179,578	2,333,484 (22.92)	17.28	28.59

NOTE: Results are for IRI data discussed previously and in detail in Bronnenberg, Kruger and Mela (2008). The data covers 313 weeks (6 years) and records weekly average prices at the store-level by product barcode.

It is perhaps more useful to examine the results at a less disaggregated. We return to the price paths for the various types of laundry detergent that we illustrated in Figure 1. We reproduce those charts but with the values indicator z included. It is clear from the figures that the model does a reasonable job of picking up the sales. Though in some cases it is not entirely clear, even from a conceptual point of view, exactly what should be categorized as a sale.

The existence of significant sales activity, as prices bounce, opens up the opportunity for shrewd consumers to "lie in wait" for sales and build inventory for later consumption. Even for relatively passive consumers the fact that the vector of prices they face may change significantly from week-to-week is likely to strongly influence purchase decisions.

Figure 3: The Sales Cycle — Laundry Detergent Prices and the Sales Indicator

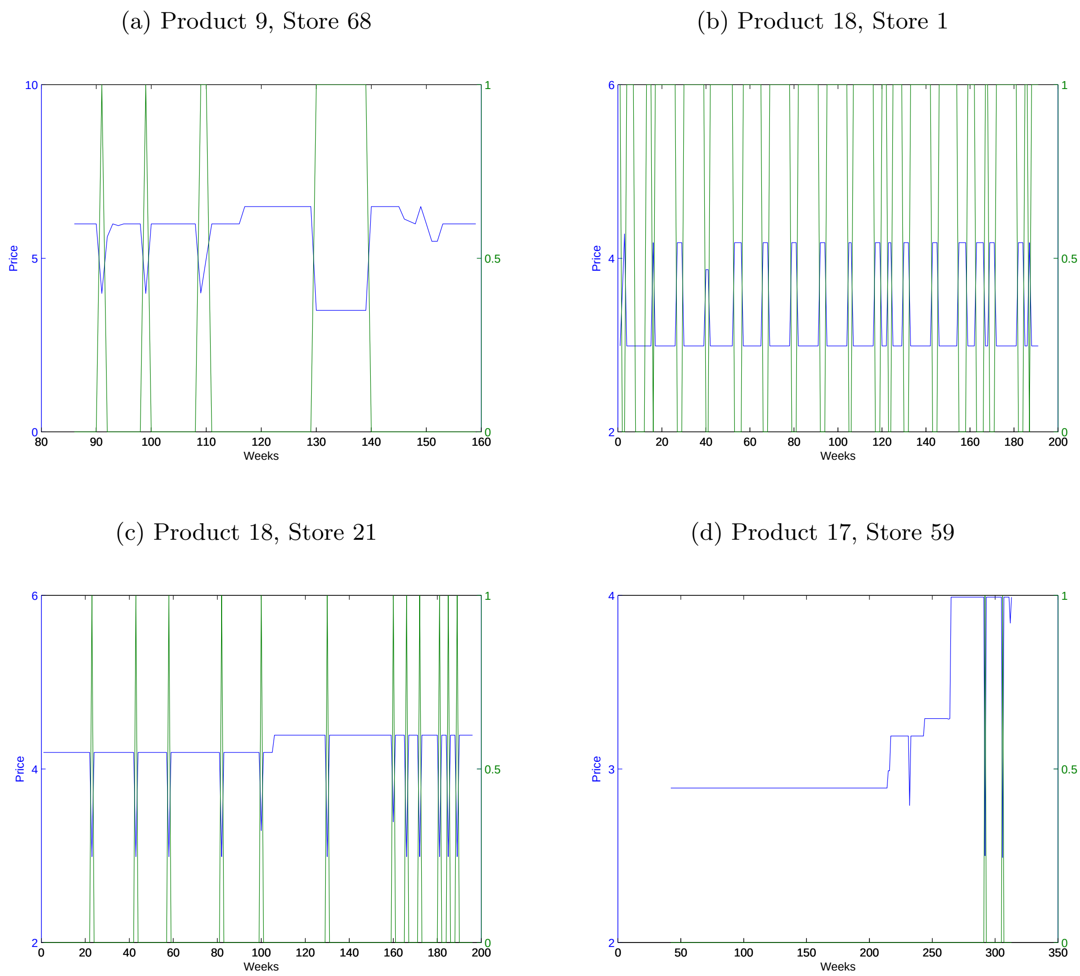

In the latter case even without stocking effects the data is likely to exhibit significant substitution effects leading to large differences between index number formulae. We return to this later. However, when some consumers do engage in stocking behaviour the substitution effects are even stronger and classic index number techniques for constructing cost-of-living indexes, which assume that purchase and consumption occur simultaneously, are somewhat undermined. In the classic case both purchase and consumption take place within a period of time. But if the time is relatively short, and goods are storable such as some of those we are examining, then this may not be the case.

It is possible to identify the existence of inventory shopping by examining the "state dependence" of purchases. If consumers are shopping only to fulfil current needs then current prices will be the driver of purchases. However, as can be shown in consumer inventory models (Hendel and Nevo, 2006a, 2006b), purchases will respond not only to prices but also to time between sales and more generally, previous prices, as this provides a proxy of the state of consumers' inventories. We propose a simple test of whether a household's unobserved inventory is influencing purchase decisions. Suppose that there is no state dependence in purchases, that is prior sales and prices do not influence current spending. Then consumers choices can be understood and explained within some classic structural economic model of decision-making, such as the CES cost function. However, if stockpiling does occur then some indicator of inventory levels will provide additional explanatory power and the classic consumption-purchase model will be undermined. This is the approach we take.

First, consider the CES cost function and the demand equations derived from it. When we take this to the data below we will just consider a CES cost function defined across items with a particular store. While there is certainly likely to be substitution between stores it is difficult to determine the extent of this given that we do not know the geographical proximity of the retailers in our data set. They may be very far away from each other, or at least some of them might be, limiting the propensity for substitution. Hence, we consider the CES cost function for a given store. The parameters of the CES cost function are η, the elasticity of substitution, and the various a_{is}, which reflect the quality or importance of the good in the cost function. Given this we have,

$$C(p_{st}, U) = \left(\sum_{i \in I_{st}} a_{is} p_{ist}^{1-\eta} \right)^{\frac{1}{1-\eta}} U \tag{4}$$

Then using Shepard's Lemma we have,

$$v_{ist} = \frac{a_{is} p_{ist}^{1-\sigma}}{\sum_{i \in I_{st}} a_{is} p_{ist}^{1-\sigma}}, \qquad v_{ist} \equiv \frac{p_{ist} q_{ist}}{\sum_{i \in I_{st}} p_{ist} q_{ist}} \tag{5}$$

Given Shepard's Lemma, we show in Appendix 7.3 that we have an estimable equation where everything is observable except the η,

$$\log \left(\frac{v_{isr}/\lambda_r}{v_{isu}/\lambda_u} \right) = (1-\sigma) \log \left(\frac{p_{isr}}{p_{isu}} / P_{ur} \right) \tag{6}$$

This gives an expression for relative demand as a function of relative prices. If this adequately represents demand, and stockpiling does not influence demand, then if we include additional variables related to inventories then these should have no explanatory power. Our approach is as follows. We want to identify the extent to which the consumer's inventory influences purchases but we do not observe inventory directly. But if consumers restock their inventory during a sale then we can be fairly certain that their inventory will be higher after the sale than before. Hence we focus on comparing two periods r and u in equation (6) where period r is the period immediately after a sale and period u is the period immediately prior to a sale. What we would expect, if inventory plays a role in determining spending, is that shares in period r will be lower than in period u after controlling for any price differences. In addition to including an intercept in equation (6) we also include effects for whether a display or feature existed in periods r or u. This gives the model below, where we have appended an additive error term,

$$\log\left(\frac{v_{isr}/\lambda_r}{v_{isu}/\lambda_u}\right) = \gamma_0 + (1-\sigma)\log\left(\frac{p_{isr}}{p_{isu}}/P_{ur}\right) + \gamma_1 \mathrm{display}_{isru} + \gamma_2 \mathrm{feature}_{isru} + e_{isru} \quad (7)$$

In implementing this approach we use the sales indicator obtained following estimation of the mixture model above. The results are shown in Table 2. As hypothesized, the intercept in equation (7) is mostly negative and statistically significant. The exceptions are the milk and mustard and ketchup product categories, which have a positive sign, though it is statistically insignificant. The estimates for the display and feature dummy variables are mainly of positive sign, though this varies and they are generally not statistically significant. The estimates for the elasticity of substitution are quite variable and somewhat troubling. In order for equation (6) to be valid the elasticity of substitution must be greater than one (Balk, 1999). This occurs in some cases but not all. Further work will focus on generating an expression where we can ensure that the elasticity of substitution is greater than one. However, overall we can take the results are strongly supportive of the hypothesis of stocking and inventory effects on consumer spending patterns.

The existence of stocking effects creates major problems for price index construction. As noted by Erdem, Imai and Keane (2003) and Hendel and Nevo (2006b) the existence of consumer stocking makes it appear—if the data were treated naively and the possibility of stocking ignored—that the short-run substitution effects are much larger than they really are. In terms of price index construction this leads to very significant differences between price index formulae. To see this let us define the Geometric Laspeyres price index (P_{bc}^{GL}), Geometric Paasche (P_{bc}^{GP}) and Törnqvist (P_{bc}^{T}) index between two periods b and c

Table 2: Results of Stockpiling Regression

Product	**Observations**	R^2	**Coefficients:**			
			Intercept	**Elasticity** (σ)	**Display**	**Feature**
Beer	250	0.0167	–0.1362**	1.2652***	0.4737	–0.8335
Carbonated Beverages	10,019	0.0067	–0.0732***	1.4242***	–0.0398	–0.0044
Coffee	7,587	0.0056	–0.0502***	0.5494***	0.0458	–0.0201
Deodorant	8,718	0.0339	–0.0350***	0.4911***	–0.0055	0.0541
Diapers	9,471	0.0475	–0.0196**	0.6477***	0.3432**	0.0039
Laundry Detergent	4,770	0.0005	–0.1061***	0.9916***	0.1792	–0.0534
Milk	8,005	0.0016	0.0016	1.4885***	0.1837	0.0015
Mustard and Ketchup	2,257	0.0116	0.0086	0.1245	0.2216	0.0469
Paper Towels	3,789	0.0617	–0.0356***	1.4993***	–0.2069	0.2553**
Peanut Butter	5,644	0.0155	–0.0910***	0.3122***	–0.0528	0.0708
Salty Snacks	5,765	0.0128	–0.0485***	1.7570***	–0.0508	–0.0014
Shampoo	9,377	0.0261	–0.0536***	0.6271***	–0.0744	0.0126
Soup	6,416	0.0087	–0.0837***	0.5901***	–0.0942	0.0061
Spaghetti Sauce	11,493	0.0019	–0.0599***	1.1128***	0.1101	0.1117
Sugar Substitute	935	0.0793	–0.0532	1.5789***	0.9133	0.23466
Toilet Tissue	5,559	0.0625	–0.0943***	1.4801***	0.0423	0.0655
Toothpaste	9,864	0.0072	–0.0400***	0.5713***	0.2121	0.0163
Yoghurt	14,983	0.0092	–0.0653***	1.6283***	0.0994	–0.0289

Note: * denotes significance at the 10% confidence level, **=5%, and ***=1%.

$$\ln P_{b,c}^{GL} = \sum_{s\in S, i\in I_s} v_{isb} \ln\left(\frac{p_{isc}}{p_{isb}}\right), \qquad v_{isb} = \frac{p_{isb}q_{isb}}{\sum_{s\in S, i\in I_s} p_{isb}q_{isb}} \tag{8}$$

$$\ln P_{b,c}^{GP} = \sum_{s\in S, i\in I_s} v_{isc} \ln\left(\frac{p_{isc}}{p_{isb}}\right), \qquad v_{isc} = \frac{p_{isc}q_{isc}}{\sum_{s\in S, i\in I_s} p_{isc}q_{isc}} \tag{9}$$

$$\ln P_{b,c}^{T} = \sum_{s\in S, i\in I_s} \frac{(v_{isb} + v_{isc})}{2} \ln\left(\frac{p_{isc}}{p_{isb}}\right) \tag{10}$$

If we go ahead and calculate the chained versions of these indexes, that is calculating the movement between periods 1 and 3 by multiplying the index between 1 and 2 with that between 2 and 3, then the results are extraordinary. They reveal an enormous divergence between the chained based-weighted Laspeyres index and the chained current-weighted Paasche indexes. In Figure 4 we show the weekly chained Paasche, Laspeyres and Törnqvist indexes for laundry detergent from 2001 to 2006.

Figure 4: Weekly Chained Indexes for Laundry Detergent from 2001 to 2006

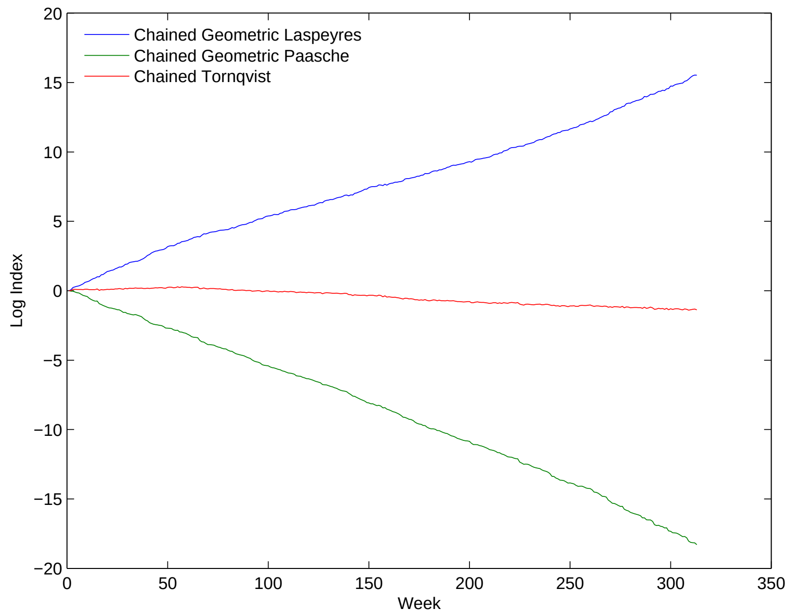

Here, in Figure 4, we have plotted the indexes on a log scale because of there extreme value. Note that a log value for the index of around 10 corresponds to an index value of around 22,000! Thus the Geometric Laspeyres and Paasche indexes are exhibiting very extreme movements. As we have argued, this reflects the sales cycle and coupled with the large change in expenditure patterns brought on by stockpiling activity. We may gain some insight into the form of the bias of these indexes by examining them in detail in light of our stochastic model of prices. Consider the decomposition of the Geometric Laspeyres index below,

$$\log P_{b,c}^{GL} = \sum_{s\in S, i\in I_s} v_{isb} \log\left(\frac{p_{isc}}{p_{isb}}\right) \tag{11}$$

$$= \sum_{s\in S, i\in I_s} v_{isb}\left[(\alpha_{isc}-\alpha_{isb}) - \beta_i(z_{isc}-z_{isb}) + (e_{isc}-e_{isb})\right] \tag{12}$$

$$= \sum_{s\in S, i\in I_s} v_{isb}(\alpha_{isc}-\alpha_{isb}) - \sum_{s\in S, i\in I_s} v_{isb}\beta_i(z_{isc}-z_{isb}) + \sum_{s\in S, i\in I_s} v_{isb}(e_{isc}-e_{isb}) \tag{13}$$

This expression is revealing. The Laspeyres index is hugely upward biased primarily because of the second term involving the sales indicator. Here the sales discount β_i is weighted by base expenditures v_{isb}. So consider the case when $z_{isc} = 1$ and $z_{isb} = 0$

so $(z_{isc} - z_{isb}) = 1$. Then the base expenditure is likely to be relatively small as it corresponds to a non-sale price. This will lead to a small negative for the whole second term, $-\beta_i(z_{isc} - z_{isb})$. But in the opposite case where $z_{isc} = 0$ and $z_{isb} = 1$, so $(z_{isc} - z_{isb}) = -1$, then the base expenditure weight is likely to be much larger and hence the whole second term, $-\beta_i(z_{isc} - z_{isb})$, can be expected to be a much larger positive number. It is this which gives the Laspeyres index its pronounced upwards bias. A similar argument holds for the Paasche index with the bias in the opposite direction.

While it is not entirely evident from the chart, because of the very large scale, even the chained Törnqvist index exhibits a high degree of drift. We provide a decomposition of the Törnqvist index below in terms of our model of prices,

$$\log P^T_{b,c} = \sum_{s \in S, i \in I_s} \frac{(v_{isc} + v_{isb})}{2} \log\left(\frac{p_{isc}}{p_{isb}}\right) \tag{14}$$

$$= \sum_{s \in S, i \in I_s} \frac{(v_{isc} + v_{isb})}{2} \Big[(\alpha_{isc} - \alpha_{isb}) - \beta_i(z_{isc} - z_{isb}) + (e_{isc} - e_{isb})\Big] \tag{15}$$

$$= \sum_{s \in S, i \in I_s} \frac{(v_{isc} + v_{isb})}{2}(\alpha_{isc} - \alpha_{isb}) - \sum_{s \in S, i \in I_s} \beta_i \frac{(v_{isc} + v_{isb})}{2}(z_{isc} - z_{isb}) + \sum_{s \in S, i \in I_s} \frac{(v_{isc} + v_{isb})}{2}(e_{isc} - e_{isb}) \tag{16}$$

What we can see is that the Törnqvist index can be decomposed into three components. The first reflects the underlying price trends in the good. The second reflects the relationship between expenditure shares and sales and the final term is a weighted average of random errors. While we would expect the last component to be essentially zero and the first to be relatively stable it is the middle term with which drift will be associated. If there is a sale in the current period but not the base period, so that $(z_{isc} - z_{isb}) = 1$, then v_{isc} is likely to be large and v_{isb} will be smaller or 'normal'. Conversely if there is a sale in the base period but not in the current period then $(z_{isc} - z_{isb}) = -1$ and we would expect v_{isc} to be large and v_{isb} to be smaller. On the face of it the fact we are using both base and current expenditure shares leads to the elimination of bias. But the problem is likely to arise because, as we argued earlier, just because prices return to normal after a sale does not mean purchases do. Inventory levels can influence purchases. Indeed, in light of the evidence presented above, in Table 2, that expenditure after a sale is smaller than expenditure before we would expect $v_{isc} + v_{isb}$ to be larger when there was a sale in period c than $v_{isc} + v_{isb}$ when there was a sale in period b. This implies that the chained

Törnqvist index will be downwardly biased as a result of this stockpiling effect. Indeed, this is mostly what we find. Consider the decomposition of the chained Törnqvist index for laundry detergents in Figure 5. Here the second component, what we call the *Beta Component*, is clearly the driver of the index and infuses it with a significant downward drift.

Figure 5: A Decomposition of the Chained Törnqvist Index for Laundry Detergent

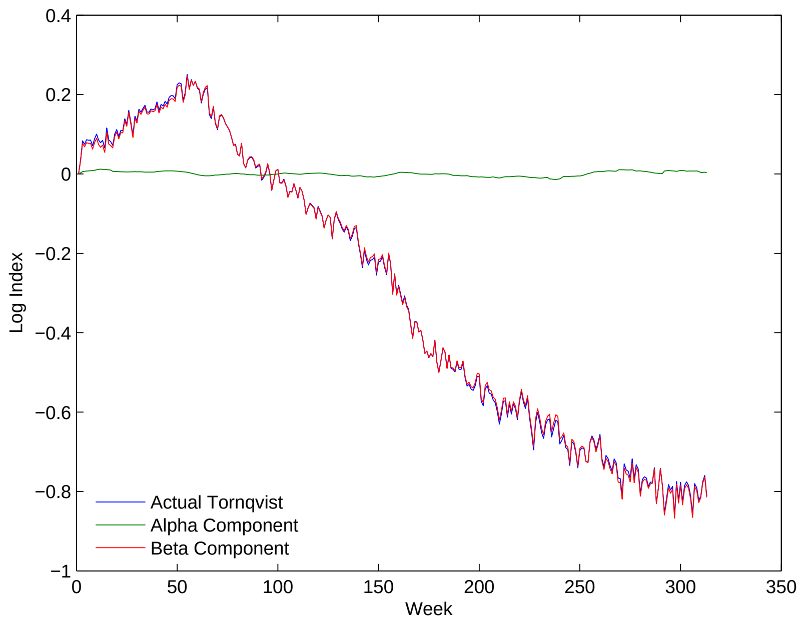

In light of these results the question arises, how do we construct a price index in the presence of such prevalent sales and stockpiling behaviour? We outline various approaches in the next section.

3 Constructing Price Indexes with Sales and Stockpiling

The economic objective in constructing a price index is often, though not always, to estimate the cost-of-living index (see National Research Council (2002) for a discussion of various approaches). This compares the ratio of cost functions with two different price vectors. At higher frequency, or in the case of semi-durable goods where the period of purchase and consumption may be quite separate, this approach can be problematic. If goods and services are not being consumed within the same period in which they are being purchased the link is broken between the observed transactions and the impact on utility. Hendel and Nevo (2006a) have previously examined the issue of consumer optimization with storable goods while Feenstra and Shapiro (2003) outlined an approach to constructing a cost-of-living index with storability. We build upon the perspective found in both these papers.

Let us define a cost minimization problem for some representative consumer which explicitly takes account of the stocking-inventory problem. Here we let x_{ist} denote the quantity of good i in store s at time t that is *consumed* and q_{ist} is the quantity *purchased*. This distinction takes account of the fact that the period of consumption may not correspond to that in which the good is purchased. Furthermore, we will consider a set of periods $t \in A$, what we term the planing horizon, over which consumption and purchase decisions are made. This may be some period like a quarter or potentially a year, in which goods that are purchased are also consumed. Of course, we acknowledge the fact that there will always be some indeterminacy at the end points. So that purchase and consumption decisions in time periods at the beginning of the planning horizon will be influenced by previous months. But we abstract from this in what follows. In order to reflect the existence of inventory, the fact that inventories may be costly to hold, and that there are some constraints on the relationship between inventories, purchases and consumption, we include an additional constraint on the consumers optimization reflected in the term $F(\mathbf{x}_{t|A}, \mathbf{q}_{t|A})$. With the addition of this term we can define the consumers cost function as the solution of the following dynamic optimization problem,

$$C(\mathbf{p}_{t|A}, \bar{U}) = \begin{cases} \text{Min.}_{\mathbf{q},\mathbf{x}} & \sum_{t \in A, s \in S_t, i \in I_{st}} p_{ist} q_{ist} \\ \text{s.t.} & U(\mathbf{x}_{t|A}) = U, \\ & F(\mathbf{x}_{t|A}, \mathbf{q}_{t|A}) = 0 \end{cases} \tag{17}$$

It is clear that the standard approach, where A contains a single time period, is contained within our 'budget horizon' cost function. Given our extended cost function a natural approach to constructing the cost-of-living index is simply to compare costs of two budget periods. That is,

$$P_{b,c|A} = \frac{C(\mathbf{p}_{t|A_c}, \bar{U})}{C(\mathbf{p}_{t|A_b}, \bar{U})} \tag{18}$$

Indeed, this is exactly what is done in many cases globally where indexes are constructed with some temporal aggregation. Where A may be a quarter or a year reflecting the highly storable nature of many goods, which can be held for many periods before being consumed. However, this creates a problem as we may be interested to estimate the cost-of-living index at a higher frequency such as monthly, weekly or even daily. It is clear that the period-to-period indexes (8)—(10) defined in the previous section are likely to bear little if any relation to (18). Given the nature of the cost function, which incorporates the potential delays in consumption, how can be estimate the cost-of-living at higher frequency?

In an interesting contribution Diewert, Fox and Ivancic (2009) have proposed a way to get around the problem of stocking and comparability by constructing a matrix of bilateral comparisons between all time periods and then multilateralizing the results by taking the geometric mean of these indexes. That is, if there are $t = 1, 2, \ldots, T$ periods for which data is available then we can calculate a full matrix of index numbers $P^T_{b,c}$ for all $b, c = 1, 2, \ldots, T$. Then the multilateralization method of Gini-Elteto-Koves-Schultz (Gini, 1931; Elteto and Koves, 1964; Szulc, 1964) can be used to transitivize these indexes. This ensures that a unique measure of price change is constructed between any two time periods. That is we have teh following GEKS indexes,

$$P^{GEKS}_{b,c} = \prod_{r \in A_T} \left(\frac{P^T_{r,c}}{P^T_{r,b}} \right)^{\frac{1}{|A_T|}} \tag{19}$$

In actual fact the approach of Diewert, Fox and Ivancic (2009) was somewhat more involved than this. They proposed to use a rolling year approach where a subset of time periods was used to undertake the multilateral comparisons and then the price levels of the two most recent periods were compared and this was used in constructing a chained index. This approach has some potential and we will explore it in the empirical section that follows. It is simple and straightforward computationally and it does not require the estimation of an econometric model of prices. However, it could be criticized as being based upon a questionable model of purchase and consumption decisions which is stretching credulity at higher frequency such as weekly or even monthly.

We suggest an alternative approach which weds the need for high frequency index numbers with the fact that the budgeting period is often inconveniently long. Consider the cost of obtaining a given level of utility, over a budget horizon A, given the particular distribution for prices that prevailed in some period t. In particular suppose that we know this distribution, or can estimate it, then we could do the following. We could generate a realization from that distribution which reflected the price level and the frequency and size of sales. Then we could use this artificially generated set of prices over a budgeting period of size $|A|$ and compare this with the base set of prices. Let $\tilde{\mathbf{p}}^c_{t|A_b}$ denote a realization from the distribution of prices in period c then we could consider the following index with regard to some reference planning horizon A_b,

$$\tilde{P}_{r,c|A_r} = \frac{C(\tilde{\mathbf{p}}^c_{t|A_r}, U)}{C(\mathbf{p}_{t|A_r}, U)} \tag{20}$$

While this approach is artificial, in that we construct a hypothetical scenario for the comparison period by trying to approximate the data generating process for prices in period c, it does have some appeal over alternative approaches. It respects the fact that consumption and purchase decisions are made over a set of periods—a budget horizon—rather than in a single period. In this regard the data for the the period of interest c is used to estimate the distribution then we create a hypothetical comparison period stretching the data over a set of periods of size $|A_r|$. While this approach has some appeal we do not have an implementable index number approach as the precise form of the cost function is unknown. There are potentially a number of ways to proceed. We could derive bounds upon the cost-of-living index outlined above or we potentially hypothesize a functional form for the cost function.

First to the second approach. Since Diewert (1976) it has been known that we may derive certain index numbers given a particular specification for the cost function. Unfortunately, this approach is unlikely to yield simple and intuitive answers. The problem is that we don't observe the decisions made by consumers with the artificial price vector $\tilde{\mathbf{p}}^c_{t|A_r}$. Hence, unlike in the usual case where we do observe the choices made by consumers under two price regimes, proceeding further is difficult. Feenstra and Shapiro (2003) ostensibly address the same issue as we have set out and they adopt an exact index number approach, using the Translog cost function and Törnqvist price index. But they do not really resolve the issue, merely compare one budgeting horizon with another. There methodology does not provide a means to obtain a high frequency measure of price change.

Second, consider equation (20). We can construct a Laspeyres upper bound on this cost-of-living index. If consumers undertake exactly the same purchasing plan facing the price draw $\tilde{\mathbf{p}}^c_{A_r}$ as they did when facing the prices over the budgeting horizon A_r

then their consumption possibilities and utility will be the same. In addition, as the previous purchase and consumption plan was feasible then it will also be feasible under the particular realization of prices that we have generated. This gives the inequality,

$$\tilde{P}_{r,c|A_r} = \frac{C(\tilde{\mathbf{p}}^c_{t|A_r}, U)}{C(\mathbf{p}_{t|A_r}, U)} \leq \frac{\sum_{t\in A_r, s\in S_t, i\in I_{st}} \tilde{p}^c_{ist} q_{ist}}{\sum_{t\in A_r, s\in S_t, i\in I_{st}} p_{ist} q_{ist}} \equiv \tilde{P}^L_{r,c|A_r} \tag{21}$$

It is clear that the Laspeyres index $\tilde{P}^L_{r,c|A_r}$ provides a transitive comparison between any two periods as we can think of them being compared relative to reference period A_r. However, this index is stochastic, in the sense that it depends upon the particular draw of prices from the distribution of period c prices. In this regard a natural approach is to focus on the expected value of the index. Clearly this abstracts from the noise introduced by the particular realization of prices. Hence we can construct,

$$\tilde{P}^{EL}_{r,c|A_r} = \frac{\mathrm{E}\left[\sum_{t\in A_r, s\in S_t, i\in I_{st}} \tilde{p}^c_{ist} q_{ist}\right]}{\sum_{t\in A_r, s\in S_t, i\in I_{st}} p_{ist} q_{ist}} \tag{22}$$

We may estimate such an expectation by analytical methods or alternatively by simulation.

Finally, it is natural in the present context, given our statistical model of prices, to consider a stochastic approach to constructing index numbers. This sits outside the cost-of-living framework to index construction but is an approach which has a long and established pedigree—according to Diewert (1995) dating back to Jevons, Edgeworth and Bowley. Another important reference is Theil (1967). It is possible to derive an index number approach, within this framework, to the problem at hand. Let us use the argument of Theil (1967) and suppose that we want to measure the difference in the cost of a unit of expenditure in one period compared with our artificial period. Then it is natural to focus on log prices, as these reflect differences, and it also natural to use expenditure shares as weights as we may think of these as the probability that a price will be selected if probabilities are drawn proportional to expenditure. If we again consider a hypothetical realization of prices from the distribution of prices in period c, i.e. $\tilde{\mathbf{p}}^c_{t|A_r}$, then this gives,

$$\log \tilde{P}^G_{r,c|A_r} = \sum_{t\in A_r, s\in S_t, i\in I_{st}} v_{ist} \log\left(\frac{\tilde{p}^c_{ist}}{p_{ist}}\right) \tag{23}$$

As will be seen, this index is particularly convenient given that we have a model for the log of prices. It is clear that if we wish to compare some time period c with another time period, say b, then the reference period r plays only a limited role. It provides the weights only. Indeed the expectation of the price index between the two periods, linked through the reference period, can be written as,

$$\log \tilde{P}^{EG}_{b,c|A_r} = \mathrm{E}\left[\log \tilde{P}^{EG}_{r,c|A_r} - \log \tilde{P}^{EG}_{r,b|A_r} \right] \tag{24}$$

$$= \sum_{t \in A_r, s \in S_t, i \in I_{st}} v_{ist|A_r} \mathrm{E}\left[\log \left(\frac{\tilde{p}^c_{ist}}{\tilde{p}^b_{ist}} \right) \right] \tag{25}$$

$$= \sum_{t \in A_r, s \in S_t, i \in I_{st}} v_{ist|A_r} \left[(\alpha_{isc} - \alpha_{isb}) - \beta_i (\omega_{ic} - \omega_{ib}) \right] \tag{26}$$

$$= \sum_{t \in A_r, s \in S_t, i \in I_{st}} v_{ist|A_r} (\alpha_{isc} - \alpha_{isb}) - \sum_{t \in A_r, s \in S_t, i \in I_{st}} v_{ist|A_r} \beta_i (\omega_{ic} - \omega_{ib}) \tag{27}$$

This index is made up of two components. The first is what might be described as a pure price index. The second represents an adjustment for the *expected* sales frequency. This means that the index reflects sales. But rather than the actual realization of sales it is the mean change in sales is recorded. This is likely to mean that the index is significantly smoother than indexes which include the sales effect directly.

3.1 A Comparison of the Various Approaches

There are some key differences between the approach embodied in the Expected Laspeyres, $\tilde{P}^{EL}_{r,c|A_r}$, and Expected Geometric, $\tilde{P}^{EG}_{r,c|A_r}$, indexes in comparison with that of Diewert, Fox and Ivancic (2009). Both provide transitive multilateral comparisons of prices yet both take quite a different perspective on this. The former approach is to create a reference period A_r, which spans a number of time periods. All other periods are compared with this reference budgeting horizon. This can be thought of as very much analogous to the average basket method in international comparisons (Hill, 1997). In contrast the Ivancic, Fox, Diewert (2009) approach suggests an alternative multilateral method which compares all individual time periods with all other periods and averages across these. This can be illustrated using Figure 6. If a node—a bold circle—represents a high frequency time period, such as a week, then we may consider two quite different approaches. The average basket approach uses what Hill (1997) refers to as a star comparison method. Here each node is compared with some reference point, the large circle. In comparison the GEKS method makes all possible bilateral comparisons, as illustrated in the second panel of Figure 6, and averages the results.

Figure 6: Multilateral Methods Compared

(a) Average Basket Method: (b) GEKS Method:

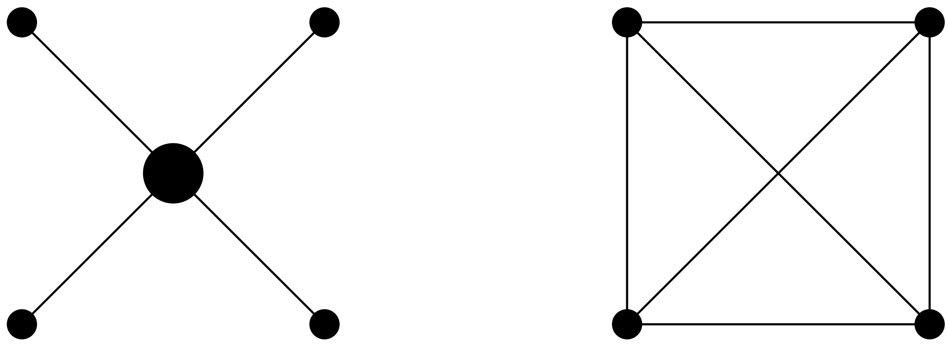

Another feature which distinguishes the two general approaches is the extent to which they satisfy the *multiperiod identity test.* If we denote $P(.)$ by the price index formula then this requires that,

$$P(\mathbf{p}_1, \mathbf{p}_2, \mathbf{q}_1, \mathbf{q}_2)P(\mathbf{p}_2, \mathbf{p}_3, \mathbf{q}_2, \mathbf{q}_3)P(\mathbf{p}_3, \mathbf{p}_1, \mathbf{q}_3, \mathbf{q}_1) = 1 \tag{28}$$

That is, suppose we have four time periods; 1, 2, 3 and 4. Furthermore, period 4 has exactly the same prices and quantities as period 1 then the chain of indexes should be equal to 1. For many indexes this will not be the case. In general the GEKS method satisfies this axiom (at least it does abstracting from the fact that the available set of items changes somewhat over time). However, satisfying this axiom may not be enough. The problem that we face in constructing price indexes at high frequency and with stockpiling is that often we may have the same prices for an item in two periods but different quantities. This may arise, as we have seen, because spending decisions at high frequency depend not just upon current prices but also previous prices and whether a good was on sale or not. A stronger requirement then might be the following,

$$P(\mathbf{p}_1, \mathbf{p}_2, \mathbf{q}_1, \mathbf{q}_2)P(\mathbf{p}_2, \mathbf{p}_3, \mathbf{q}_2, \mathbf{q}_3)P(\mathbf{p}_3, \mathbf{p}_1, \mathbf{q}_3, \mathbf{q}_4) = 1 \tag{29}$$

This requirement says that, even if quantities are different in period 4 from period 1, with prices the same, then the indexes should multiply to one. While there is no guarantee that the GEKS method satisfy this requirement the average basket method does by construction. Because in constructing the price index only the quantities or expenditure weights of the reference period are used. The quantities in the periods being compared play no role. This is potentially quite important.

We have already discussed the form of the chained Geometric Laspeyres, Paasche and Törnqvist indexes and argued that they may be biased in the presence of stocking behaviour. Let us now examine the multilateral Törnqvist index, which adopts the transitivisation approach of Elteto and Koves (1964) and Szulc (1964), and has been recently suggested in the scanner data context by Ivancic, Fox and Diewert (2009). That is, again assuming a constant set of products and stores across time, the temporal price parities are constructed as,

$$\log P_{b,c}^{GEKS} = \frac{1}{|A_T|} \sum_{a \in A_T} \left(\log P_{ac}^{T} - \log P_{ab}^{T} \right) \tag{30}$$

$$= \frac{1}{|A_T|} \sum_{a \in A_T} \left(\sum_{s \in S, i \in I_s} \frac{(v_{isa} + v_{isc})}{2} \log \left(\frac{p_{isc}}{p_{isa}} \right) - \sum_{s \in S, i \in I_s} \frac{(v_{isa} + v_{isb})}{2} \log \left(\frac{p_{isb}}{p_{isa}} \right) \right) \tag{31}$$

$$= \frac{1}{|A_T|} \sum_{a \in A_T} \left(\sum_{s \in S, i \in I_s} \frac{(v_{isc} + v_{isa})}{2} \left[(\alpha_{isc} - \alpha_{isa}) - \beta_i (z_{isc} - z_{isa}) + (e_{isc} - e_{isa}) \right] \right.$$
$$\left. - \sum_{s \in S, i \in I_s} \frac{(v_{isb} + v_{isa})}{2} \left[(\alpha_{ib} - \alpha_{ia}) - \beta_i (z_{isb} - z_{isa}) + (e_{isb} - e_{isa}) \right] \right) \tag{32}$$

$$= \frac{1}{|A_T|} \sum_{a \in A_T} \sum_{s \in S, i \in I_s} \left[\frac{(v_{isc} + v_{isa})}{2} (\alpha_{isc} - \alpha_{isa}) - \frac{(v_{isb} + v_{isa})}{2} (\alpha_{isb} - \alpha_{isa}) \right]$$
$$- \frac{1}{|A_T|} \sum_{a \in A_T} \sum_{s \in S, i \in I_s} \beta_i \left[\frac{(v_{isc} + v_{isa})}{2} (z_{isc} - z_{isa}) - \frac{(v_{isb} + v_{isa})}{2} (z_{isb} - z_{isa}) \right]$$
$$+ \frac{1}{|A_T|} \sum_{a \in A_T} \sum_{s \in S, i \in I_s} \left[\frac{(v_{isc} + v_{isa})}{2} (e_{isc} - e_{isa}) - \frac{(v_{isb} + v_{isa})}{2} (e_{isb} - e_{isa}) \right] \tag{33}$$

We can write each of these components in such a way that we have,

$$\log P_{b,c}^{GEKS} = \sum_{s \in S, i \in I_s} \left[\frac{(v_{isc} + \bar{v}_{is})}{2} (\alpha_{isc} - \bar{\alpha}_{is}) - \frac{(v_{isb} + \bar{v}_{is})}{2} (\alpha_{isb} - \bar{\alpha}_{is}) \right]$$
$$- \sum_{s \in S, i \in I_s} \beta_i \left[\frac{(v_{isc} + \bar{v}_{is})}{2} (z_{isc} - \bar{z}_{is}) - \frac{(v_{isb} + \bar{v}_{is})}{2} (z_{isb} - \bar{z}_{is}) \right]$$

$$+ \sum_{s\in S, i\in I_s} \left[\frac{(v_{isc} + \bar{v}_{is})}{2}(e_{isc} - \bar{e}_{is}) - \frac{(v_{isb} + \bar{v}_{is})}{2}(e_{isb} - \bar{e}_{is}) \right], \quad \bar{x}_{is} = \frac{1}{|A_T|} \sum_{a\in A_T} x_{isa} \tag{34}$$

This shows that the multilateral Törnqvist index is equal to a pure price index, a sales index and some random errors. The sales index is constructed as the weighted average of the difference between the sales indicators. There are certain similarities between this index and that which we suggested earlier. However, there is one major difference. This is that in the multilateral Törnqvist index price change could be recorded even when no prices have changed. To see this note that if $z_{isc} = z_{isb}$ and $\alpha_{isc} = \alpha_{isb}$ for all i and s in (34) then unless $v_{isc} = v_{isb}$ price change will be recorded. This is the major disadvantage of this index in comparison with our suggested indexes and is related to the fact that the GEKS method uses the relevant period's expenditure weights in constructing comparisons rather than fixing them. It maximises characteristicity (Drechsler, 1973) at the expense of allowing weighting to change and reflecting spurious price changes. In contrast the Expected Laspeyres and Expected Geometric indexes, because they use constant expenditure or quantity weights, will only record price change when prices do actually differ.

With an understanding of how the various price indexes differ conceptually we now turn to the empirical magnitudes by constructing these price indexes using the IRI scanner data. This enables us to get a clearer handle on the empirical magnitudes of bias and volatility and provides some guidance in practical implementation of the approach.

4 An Empirical Examination of the Price Indexes

In the preceding sections we have outlined an index number methodology which is likely to be useful in constructing high frequency price indexes using scanner data. Let us now apply this approach to the IRI data.

We make use of a very large scanner data set made publicly available by IRI (see Bronnenberg, Kruger and Mela, 2008). The data includes weekly sales data for a number of products over a number of markets (regions) for the years 2001 to 2006. The data comes from a variety of supermarket and drug stores across a range of chains which were supplying data to IRI over the period. The breadth and depth of the data provide a ideal basis on which to test methods for constructing high frequency price indexes. In our application we use just a small portion of the data and consider just a few of the product categories in detail. These are; carbonated beverages, coffee, laundry detergent, milk, mustard and ketchup, peanut butter, shampoo and toothpaste. We make use of all 6 years of the data. But in constructing our data we remove items which are sold relatively infrequently—with fewer than 52 observations over 6 years—and then randomly sample across stores and items. This subsampling of the data is necessary at this preliminary stage of the investigation because the IRI data is very large, containing many millions of observations. Use of all of it would impose considerable computational burdens.

In summarizing the results we focus primarily on the multilateral methods outlined previously. In almost all cases the chained Laspeyres, Paasche and Törnqvist indexes exhibit very significant bias. We illustrate by including the chained Törnqvist index in the first column of Figures 7 and 8. In most cases, barring that of milk and salty snacks, the chained Törnqvist index is significantly downwardly biased. This reflects the arguments made in section 2, equation (16), regarding the correlation between the sales indicator and expenditure shares. Given this, it is somewhat surprising that not all product categories exhibit a chained Törnqvist index which is downwardly biased. The fact that milk and salty snacks do not exhibit such a trend may be the result of the particularities of the products. Neither are highly storable, milk in particular, which is likely to reduce the size of the correlation between sales and expenditure shares discussed earlier. Nevertheless, it is clear that in general the chained Törnqvist index exhibits significant drift and its use cannot be recommended.

In comparison the other indexes—the GEKS Törnqvist, Expected Laspeyres and Expected Geometric—perform relatively better. Here, given that the emphasis is on developing methods which can be implemented in real time, we have constructed each of these indexes using the rolling-year approach of Diewert, Fox and Ivancic (2009).

Let us consider the rolling year GEKS Törnqvist index shown in the first column of Figures 9 and 10. Let us focus primarily on the case of laundry detergents, Figures 9e and 9f, as it provides insights which are broadly consistent with other product categories. First, there is clearly very little inflation or deflation in prices over this period. Given this, what is most noticeable about the GEKS Törnqvist index is its volatility. The mean absolute price change per week is 2.64%. This is very significant when the trend in the index is essentially zero. This volatility is primarily driven by the sales cycle as was discussed with regard to equation (34). The GEKS Törnqvist index, while it averages price change over many periods, includes the effect of sales directly in the index. This is not the case for either the Expected Laspeyres or Expected Geometric indexes which focus on changes in the rate of sales over time. The different in volatility between the indexes is evident in Figure 9e.

Given that the GEKS Törnqvist index includes the effects of sales directly into the index there is a strong relationship between it and the probability-of-sale indicator. If we continue to focus on laundry detergents then it can be seen in the second period, in Figure 7f, there is a marked fall in the incidence of sales. This leads to a corresponding jump in the GEKS Törnqvist index. This is much less apparent in either of the other indexes in Figure 9e. The result of this is that the GEKS Törnqvist index has a variance which is many orders of magnitude higher than that of either the Expected Laspeyres or Expected Geometric indexes.

However, comparing the variance of the Expected Laspeyres or Expected Geometric indexes with that of the GEKS Törnqvist index is somewhat misleading. In the latter we include 'raw' unmodeled prices while in the former we have implicitly smoothed them out through the modeling process. A less loaded comparison between the index methods is to compare their variance when we use actual prices in constructing the Expected Laspeyres and Expected Geometric indexes. Here we sample from the actual distribution of prices rather than using the parameters of the model in equations (22) and (23) to calculate the expectation. The resulting indexes are shown in the second column of Figures 9 and 10. While the Expected Laspeyres and Expected Geometric indexes are more volatile, as would be expected, there is still significantly less variability in them than for the GEKS Törnqvist index. For laundry detergent the Expected Laspeyres variance is 37.62% of that of the GEKS Törnqvist index while the Expected Geometric has just 7.28% of the variance of the GEKS Törnqvist index. This represents a significant improvement in efficiency. Put another way, the average absolute weekly change in the GEKS Törnqvist index is 2.64% while the equivalent value for the Expected Laspeyres is 1.54% and for the Expected Geometric is just 0.65%. The variance of the two average basket indexes, relative to the GEKS Törnqvist index, is shown in Table 3 for laundry detergent and the various other product categories.

With regard to bias, while all three of our preferred indexes differ, it would be a stretch to regard any of them as necessarily significantly biased. Though as we saw above, the Expected Laspeyres index can be thought of as an upper bound on the 'budget horizon' cost of living index. As the geometric mean lies below the arithmetic mean the Expected

Geometric index may give a closer approximation to the actual cost of living index. With regard to the GEKS Törnqvist index, as we have argued this may record price change even when all prices are the same. This is inherent to the methodology and reflects the fact that it uses different weighting patterns when comparing two periods. However, while this arguably results in an erroneous measure of price change it is unlikely to result in any systematic bias in either direction.

Table 3: Variance Relative to GEKS Törnqvist

Product	**Expected Laspeyres**	**Expected Geometric**
Coffee	22.12	8.60
Carbonated Beverages	27.29	10.38
Laundry Detergent	37.62	7.28
Milk	39.68	20.74
Mustard and Ketchup	5.46	2.25
Peanut Butter	14.08	7.24
Shampoo	15.33	6.19
Toothpaste	20.06	8.81

5 Conclusion

This paper has been primarily exploratory, and somewhat speculative, in nature. Its purpose has been to try to understand some of the problems with statistical agencies' use of scanner data and propose some novel solutions. Currently a very good method, the rolling year GEKS approach, is available for calculating price indexes at high frequency. This article proposed a few more approaches, which are conceptually somewhat different, which may in time be regarded as alternatives to what is the current 'industry standard'.

One contribution of the paper was to outline a method for endogenously identifying sales. This potentially has widespread use, not only in price index construction but also in marketing and elsewhere. The mixture model approach represents a very natural framework for thinking about sales and sales behaviour. There are numerous possible extensions to the model outlined above including making the sales indicator spatially and temporally dependent.

The indexes proposed in this paper are conceptually related to the average basket approach in multilateral comparisons. This kind of approach seems both conceptually sound and provides indexes which are readily calculable. The problem is that consumers do not purchase for consumption at a weekly or even a monthly frequency. They often purchase, store and consume at a later date. The purchase and consumption decisions can be very displaced in time if there are significant discounts available for some products. Our approach to this challenging problem is the following. We suppose that purchase and consumption decisions being made over some 'budget horizon'. This is a period which is as short as possible but long enough that we can regard purchase and consumption decisions as being reasonably self-contained within it. Something of the order of a year seems reasonable and this is what is used in our empirical application. We then said that we would estimate the parameters of the distribution of prices for a given period t, a short period of time much less than a budgeting horizon, and essentially create an artificial comparison period by taking simulating prices from this distribution. What we are doing is comparing an artificially created budgeting period, with prices from the distribution in period t, with an actual budgeting period. The weights from the actual budgeting period are used to make the comparison. This yields what are essentially Laspeyres-type fixed based index. This has the advantages that is conceptually reasonable, is broadly consistent with current statistical agency practice and will also ensure that the index does not change when prices themselves do not change.

This approach may provide a basis for further research on different methods for constructing prices indexes using scanner data. Up until now the 'take up' of scanner data by statistical agencies has been very disappointing. It is hoped that this paper has provided some additional tools and alternative methods which will help statistical agencies see their way through the current fog and eventually assist in them integrating scanner data into their CPIs.

6 References

Balk, B. (1999), "On Curing the CPI's Substitution and New Goods Bias," Paper Presented at the Fifth Meeting of the Ottawa Group, Reykjavik, Iceland.

BLS (2009), "Table 1 (2005-2006 Weights): Relative importance of components in the Consumer Price Indexes: U.S. city average, December 2009", Viewed 1 April 2010, www.bls.gov/cpi/cpiri05-06 2009.pdf.

Boskin, M. J., E. R. Dulberger, R. J. Gordon, Z. Griliches and D. W. Jorgenson (1997), “The CPI Commission: Findings and Recommendations”, *The American Economic Review* 87(2), 78-83.

Bronnenberg, B. J., M. W. Kruger and C. F. Mela (2008), “The IRI Marketing Data Set,” *Marketing Science* 27(4), 745748.

Deltas, G. and E. Zacharias (2004), “Sampling Frequency and the Comparison between Matched-Model and Hedonic Regression Price Indexes,” *Journal of Business and Economic Statistics* 22.

Dempster, A. P., N. M. Laird and D. B. Rubin (1977), “Maximum Likelihood from Incomplete Data via the EM Algorithm”, *Journal of the Royal Statistical Society, Series B (Methodological)* 39(1), 138

Diewert (1995), “On the Stochastic Approach to Index Numbers,” mimeo.

Diewert, E., K. J. Fox and L. Ivancic (2009), “Scanner Data, Time Aggregation and the Construction of Price Indexes,” UBC Discussion Paper 2009-48, UBC Department of Economics.

Drechler, L. (1973), “Weighting of Index Numbers in International Comparisons,” *Review of Income and Wealth* 19(1), 17-47.

Elteto, O. and P. Koves (1964), “On an Index Number Computation Problem in International Comparison”, *Statisztikai Szemle* 42, pp. 507518 (in Czech).

Erdem, T., S. Imai and M. P. Keane (2003), “Brand and Quantity Choice Dynamics Under Price Uncertainty,” *Quantitative Marketing and Economics* 1, 5-64.

Feenstra, R. C. (1994), “New Product Varieties and the Measurement of International Prices,” *American Economic Review* 84(1), 157-77.

Feenstra, R. C. and M. D. Shapiro (eds.) (2003), *Scanner Data and Price Indexes*, University of Chicago Press, US.

Feenstra, R. C. and M. D. Shapiro (2003), “High-Frequency Substitution and the Measurement of Price Indexes,” in *Scanner Data and Price Indexes* Feenstra, R. C. and M. D. Shapiro (eds.), University of Chicago Press, US.

Gini, C. (1931), “On the Circular Test of Index Numbers,” *International Review of Statistics* 9(2), 3-25.

Goldberger, A. (1968), "The Interpretation and Estimation of Cobb-Douglas Functions," *Econometrica* 36(3/4), 464-472.

Griffith, R., E. Leibtag, A. Leicester, and A. Nevo (2008), "Timing and Quantity of Consumer Purchases and the Consumer Price Index," NBER Working Paper Series, Working Paper 14433.

Haan, J. d. and H. v. d. Grient (2009), "Eliminating Chain Drift in Price Indexes Based on Scanner Data", mimeo.

Helsen, K. and D. C. Schmittlein (1992), "Some Characterizations of Stockpiling Behavior Under Uncertainty", *Marketing Letters* 3(1), 5-17.

Hendel, I. and A. Nevo (2006a), "Sales and Consumer Inventory", *The RAND Journal of Economics* 37(3), 543-561.

Hendel, I. and A. Nevo (2006b), "Measuring the Implications of Sales and Consumer Inventory Behavior", *Econometrica* 74(6), 1637-1673.

Hess, J. D. and E. Gerstner (1987), "Loss Leader Pricing and Rain Check Policy," *Marketing Science* 6(4), 358-374

Hill, R. J. (1997), "A Taxonomy of Multilateral Methods for Making International Comparisons of Prices and Quantities," *Review of Income and Wealth* 43(1), 49-69.

Kennedy, P. E. (1981), "Estimation with correctly interpreted dummy variables in semilogarithmic equations," *American Economic Review* 71(4), 801.

Melser, D. (2006), "Accounting for the Effects of New and Disappearing Goods Using Scanner Data," *Review of Income and Wealth* 52(4), 547-568.

National Research Council (2002), *At What Price? Conceptualizing and Measuring Cost-of-Living and Price Indexes*, Panel on Conceptual, Measurement, and Other Statistical Issues in Developing Cost-of-Living Indexes, Charles L. Schultze and Christopher Mackie (eds.), Committee on National Statistics, Division of Behavioral and Social Sciences and Education, Washington, DC: National Academy Press.

Reinsdorf, M. (1999), "Using Scanner Data to Construct CPI Basic Component Indexes," *Journal of Business and Economic Statistics* 17(2), 152160.

Sato, K. (1976), "The Ideal Log-Change Index Number," *Review of Economics and Statistics* 58(2), 223-28.

Silver, M. (1995), "Elementary Aggregates, Micro-Indices and Scanner Data: Some Issues in the Compilation of Consumer Price Indices," *Review of Income and Wealth* 41(4), 427438.

Silver, M. (1999), "An Evaluation of the Use of Hedonic Regressions for Basic Components of Consumer Price Indices," *Review of Income and Wealth* 45(1), 4156.

Silver, Mick and B. Webb (2002), "The Measurement of Inflation: Aggregation at the Basic Level, *Journal of Economic and Social Measurement* 28(12), 2135.

Silver, M. and S. Heravi (2001), "Scanner Data and the Measurement of Inflation," *The Economic Journal* 111(June), 383-404.

Statistics Netherlands (2003), *Press Release*, Viewed 27 April 2011, www.cbs.nl/NR/rdonlyres/A703BC78-2345-4E39-B16A-11F1B64BD171/0/pb03e165.pdf

Szulc, B. (1964), "Indices for Multiregional Comparisons", *Przeglad Statystyczny* 3, 239-254.

Theil, H. (1967), *Economics and Information Theory*, Amsterdam: North-Holland.

Triplett, J. E. (2003), "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*, University of Chicago Press, US, 15164.

Van Heerde, H. J., P. S. H. Leeflang and D. R. Wittink (2000), "The Estimation of Pre- and Postpromotion Dips with Store-Level Scanner Data," *Journal of Marketing Research* 37(3), 383-395.

Varian, H. (1980), "A model of sales," *The American Economic Review* 70(4), 651-9.

Vartia, Y. O. (1976), "Ideal Log-Change Index Numbers," *Scandinavian Journal of Statistics* 3, 121-26.

7 Appendix

7.1 Estimation of the Mixture Model

Let us define the set of unknown parameters as $\Omega = \{\boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\sigma}^2, \boldsymbol{\omega}\}$ and some estimate of the parameters at iteration r as $\Omega_{(r)} = \{\boldsymbol{\alpha}_{(r)}, \boldsymbol{\beta}_{(r)}, \boldsymbol{\sigma}^2_{(r)}, \boldsymbol{\omega}_{(r)}\}$. The EM algorithm has two steps (Ripley, 2004):

$$\text{E}-\text{step}: \qquad \text{Find } \; Q(\Omega,\Omega_{(r)}) = E_{z_{ist}|\Omega_{(r)}}\left[\sum_{t\in A_T,s\in S,i\in I_{st}} w_{ist}\log p(y_{ist},z_{ist}|\Omega)\right] \tag{35}$$

$$\text{M}-\text{step}: \qquad \text{Choose } \; \Omega_{(r+1)} \; \text{so as to maximise} \; Q(\Omega,\Omega_{(r)}). \tag{36}$$

Here $p(y_{ist}, z_{ist}|\Omega)$ is the probability density of the observed data y_{ist} and the unobserved 'data' z_{ist} given the parameters Ω. Note that in addition to the standard approach we have introduced some weights w_{ist} to the likelihood which could be used to reflect the economic importance of the observation.

To demonstrate the convergence of the EM algorithm following Ripley (2004, p.335) we use the identity, $p(y_{ist}|\Omega) = \frac{p(y_{ist},z_{ist}|\Omega)}{p(z_{ist}|y_{ist},\Omega)}$, take logs, and then take the expectation treating z_{ist} as a random variable conditional on some estimate of the parameters $\Omega_{(r)}$, which is what we mean by $E_{z_{ist}|\Omega_{(r)}}[.]$. Furthermore, noting that the (weighted) log likelihood $\log L(\Omega|\mathbf{y}) = \sum_{t\in A_T,s\in S,i\in I_{st}} w_{ist}\log p(y_{ist}|\Omega)$, yields,

$$\log L(\Omega|\mathbf{y}) = Q(\Omega,\Omega_{(r)}) - E_{z_{ist}|\Omega_{(r)}}\left[\sum_{t\in A_T,s\in S,i\in I_{st}} w_{ist}\log p(z_{ist}|y_{ist},\Omega)\right] \tag{37}$$

Considering the second term note that as, $\log x \leq x-1$, we have,

$$E_{z_{ist}|\Omega_{(r)}}\left[\log\left(\frac{p(z_{ist}|y_{ist},\Omega)}{p(z_{ist}|y_{ist},\Omega_{(r)})}\right)\right] \leq \qquad E_{z_{ist}|\Omega_{(r)}}\left[\frac{p(z_{ist}|y_{ist},\Omega)}{p(z_{ist}|y_{ist},\Omega_{(r)})}\right] - 1 = 0 \tag{38}$$

Therefore,

$$E_{z_{ist}|\Omega_{(r)}}\left[\log p(z_{ist}|y_{ist},\Omega)\right] \leq E_{z_{ist}|\Omega_{(r)}}\left[\log p(z_{ist}|y_{ist},\Omega_{(r)})\right] \tag{39}$$

Hence the second term in (37) is maximized at $\Omega_{(r)}$. Any other value, say $\Omega_{(r+1)}$ will lead to a lower value for this term by (39) and, because of the negative sign in front of it in (37), a higher value for the log likelihood function. Hence choosing $\Omega_{(r+1)}$, such that $Q(\Omega_{(r+1)},\Omega_{(r)}) > Q(\Omega_{(r)},\Omega_{(r)})$, must raise the value of the log likelihood.

Turning now to the steps of the EM algorithm and the identification of the elements of $Q(\Omega,\Omega_{(r)})$. For our model (1) and (2) above we can write $Q(\Omega,\Omega_{(r)})$ as,

$$Q(\Omega,\Omega_{(r)}) = E_{z_{ist}|\Omega_{(r)}}\left[\sum_{t\in A_T, s\in S, i\in I_{st}} w_{ist}\log p(y_{ist}, z_{ist}|\Omega)\right] \tag{40}$$

$$\begin{aligned} &= E_{z_{ist}|\Omega_{(r)}}\left[\sum_{t\in A_T, s\in S, i\in I_{st}} w_{ist}\left\{\frac{(1-z_{ist})}{2}\left(2\log(1-\omega_i)-\log(2\pi\sigma_i^2)-\left(\frac{y_{ist}-\alpha_{ist}}{\sigma_i}\right)^2\right)\right.\right. \\ &\left.\left. +\frac{z_{ist}}{2}\left(2\log(\omega_i)-\log(2\pi\sigma_i^2)-\left(\frac{y_{ist}-(\alpha_{ist}-\beta_i)}{\sigma_i}\right)^2\right)\right\}\right] \end{aligned} \tag{41}$$

Now to the maximization of $Q(\Omega,\Omega_{(r)})$. We can use, $E[z_{ist}=1|\Omega_{(r)}] = p(z_{ist}=1|\Omega_{(r)})$, and Bayes rule, $p(z_{ist}=1|\Omega_{(r)}) = \frac{p(z_{ist}=1,\Omega_{(r)})}{p(\Omega_{(r)})}$ to estimate the following,

$$p(z_{ist}=1|\Omega_{(r)}) = \frac{\omega_{i(r)}\pi(y_{ist}|\alpha_{i(r)}-\beta_{i(r)},\sigma^2_{i(r)})}{(1-\omega_{i(r)})\pi(y_{ist}|\alpha_{i(r)},\sigma^2_{i(r)})+\omega_{i(r)}\pi(y_{ist}|\alpha_{i(r)}-\beta_{i(r)},\sigma^2_{i(r)})} \tag{42}$$

Where $\pi(x|\mu,v^2)$ denotes the probability density function for variable x from the normal distribution with mean μ and variance v^2. To estimate the remaining parameters using our conditional expectation of z_{ist}, that is $z_{ist(r+1)}$, we substitute these estimates into $Q(\Omega|\Omega_{(r)})$ and maximize with respect to the unknown parameters. In order to estimate ω_{it} we take the mean of the $z_{ist(r+1)}$,

$$\omega_{it(r+1)} = \sum_{s\in S}\left(\frac{w_{ist}}{\sum_{s\in S} w_{ist}}\right) z_{ist(r+1)} \tag{43}$$

For the mean parameters, taking the variance $\sigma^2_{i(r)}$ as given, we minimize the sum-of-squares,

$$S = \sum_{t\in A_T, s\in S, i\in I_{st}} w_{ist}\left((1-z_{ist(r+1)})\left(\frac{y_{ist}-\alpha_{ist}}{\sigma_{i(r)}}\right)^2 + z_{ist(r+1)}\left(\frac{y_{ist}-(\alpha_{ist}-\beta_i)}{\sigma_{i(r)}}\right)^2\right) \tag{44}$$

Finally, given the other parameter estimates, the product specific variances can be estimated as,

$$\begin{aligned}\sigma^2_{i(r+1)} = \sum_{t\in A_T, s\in S}\left(\frac{w_{ist}}{\sum_{t\in A_T, s\in S} w_{ist}}\right)&\Big[(1-z_{ist})\left(y_{ist}-\alpha_{it(r+1)}\right)^2 \\ &+ z_{ist}\left(y_{ist}-\alpha_{ist(r+1)}+\beta_{i(r+1)}\right)^2\Big]\end{aligned} \tag{45}$$

The two steps of the EM algorithm can be iterated until convergence is achieved.

7.2 Variance of the Mean for a Mixture Distribution

Here we derive the variance of an observation from the mixture distribution,

$$
\begin{aligned}
\mathrm{Var}(\ln p_{ist}) &= \mathrm{E}\left[(\ln p_{ist})^2\right] - \mathrm{E}\left[\ln p_{ist}\right]^2 && (46)\\
&= \left(\omega_{it}\left((\alpha_{ist}-\beta_i)^2+\sigma_i^2\right)+(1-\omega_{it})\left((\alpha_{ist})^2+\sigma_i^2\right)\right) \\
&\qquad - \left(\omega_{it}\left(\alpha_{ist}-\beta_i\right)+(1-\omega_{it})\left(\alpha_{ist}\right)\right)^2 && (47)\\
&= \sigma_i^2+\alpha_{ist}^2+\omega_{it}\beta_i^2-2\omega_{it}\alpha_{ist}\beta_i-\left(\alpha_{ist}^2+\omega_{it}^2\beta_i^2-2\omega_{it}\alpha_{ist}\beta_i\right) && (48)\\
&= \sigma_i^2+\omega_{it}(1-\omega_{it})\beta_i^2 && (49)
\end{aligned}
$$

7.3 Derivation of CES Demand Function

If we start with Shepard's Lemma and minus off the equation for the same product in a different time period. This gives,

$$
\log\left(\frac{v_{isr}}{v_{isu}}\right) = (1-\sigma)\log\left(\frac{p_{isr}}{p_{isu}}\right) - \log\left(\frac{\sum_{s\in S,i\in I_{sr}} a_{is}p_{isr}^{1-\sigma}}{\sum_{s\in S,i\in I_{sr}} a_{is}p_{isu}^{1-\sigma}}\right) \tag{50}
$$

Now, using Shepard's Lemma, the expression on the far right hand side of this equation can be written as,

$$
\begin{aligned}
\left(\frac{\sum_{s\in S,i\in I_{sr}} a_{is}p_{isr}^{1-\sigma}}{\sum_{s\in S,i\in I_{su}} a_{is}p_{isu}^{1-\sigma}}\right) &= \frac{p_{isr}^{1-\sigma}/v_{isr}}{p_{isu}^{1-\sigma}/v_{isu}},\ \forall s=1,2,\ldots,S,\ i\in I_{sr,su},\ I_{sr,su}=I_{sr}\cap I_{su} && (51)\\
&= \frac{v_{isu}}{v_{isr}}\frac{p_{isr}^{1-\sigma}}{p_{isu}^{1-\sigma}},\ \forall s=1,2,\ldots,S,\ i\in I_{sr,su} && (52)\\
&= \left(\frac{v_{isu}}{v_{isr}}\right)\left(\frac{p_{isr}}{p_{isu}}\right)^{1-\sigma},\ \forall s=1,2,\ldots,S,\ i\in I_{sr,su} && (53)\\
&= \left(\frac{\lambda_u}{\lambda_r}\right)\left(\frac{\left(\frac{v_{isu}}{\lambda_u}\right)}{\left(\frac{v_{isr}}{\lambda_r}\right)}\right)\left(\frac{p_{ir}}{p_{iu}}\right)^{1-\sigma},\ \forall s=1,2,\ldots,S,\ i\in I_{sr,su}, \\
&\qquad \lambda_r=\sum_{s\in S,i\in I_{sr,su}} v_{isr},\ \lambda_u=\sum_{s\in S,i\in I_{sr,su}} v_{isu} && (54)\\
&= \left(\frac{\lambda_u}{\lambda_r}\right)\left(\frac{\bar{v}_{iu}}{\bar{v}_{ir}}\right)\left(\frac{p_{ir}}{p_{iu}}\right)^{1-\sigma},\ \forall s=1,2,\ldots,S,\ i\in I_{sr,su}, \\
&\qquad \bar{v}_{isr}=\frac{s_{isr}}{\lambda_r},\ \bar{v}_{isu}=\frac{s_{isu}}{\lambda_u} && (55)\\
&= \left(\frac{\lambda_u}{\lambda_r}\right)\prod_{s\in S,i\in I_{sr,su}}\left[\left(\frac{\bar{v}_{isu}}{\bar{v}_{isr}}\right)^{w_{i|I_{sr,su}}}\left(\frac{p_{isr}}{p_{isu}}\right)^{\frac{w_{i|I_{sr,su}}}{1-\sigma}}\right],
\end{aligned}
$$

$$w_{i|I_{sr,su}} = \frac{\left(\frac{\bar{v}_{isr}-\bar{v}_{isu}}{\ln \bar{v}_{isr}-\ln \bar{v}_{isu}}\right)}{\sum_{s\in S,i\in I_{sr,su}}\left(\frac{\bar{v}_{isr}-\bar{v}_{isu}}{\ln \bar{v}_{isr}-\ln \bar{v}_{isu}}\right)} \tag{56}$$

$$= \left(\frac{\lambda_u}{\lambda_r}\right)\left(\prod_{i\in I_{ru}}\left(\frac{p_{isr}}{p_{isu}}\right)^{\frac{w_{i|I_{sr,su}}}{1-\sigma}}\right) \tag{57}$$

$$= \left(\frac{\lambda_u}{\lambda_r}\right) P_{ur}^{1-\sigma} \tag{58}$$

Where we have used the fact that,

$$\ln \prod_{s\in S,i\in I_{sr,su}}\left(\frac{\bar{v}_{isu}}{\bar{v}_{isr}}\right)^{w_{i|I_{sr,su}}} = \sum_{s\in S,i\in I_{sr,su}} w_{i|I_{sr,su}}(\ln \bar{v}_{isu} - \ln \bar{v}_{isr}) \tag{59}$$

$$= \sum_{s\in S,i\in I_{sr,su}} \frac{\left(\frac{\bar{v}_{isr}-\bar{v}_{isu}}{\ln \bar{v}_{isr}-\ln \bar{v}_{isu}}\right)}{\sum_{s\in S,i\in I_{sr,su}}\left(\frac{\bar{v}_{isr}-\bar{v}_{isu}}{\ln \bar{v}_{isr}-\ln \bar{v}_{isu}}\right)}(\ln \bar{v}_{isu} - \ln \bar{v}_{isr}) \tag{60}$$

$$= -\sum_{s\in S,i\in I_{sr,su}} \frac{\bar{v}_{isu}-\bar{v}_{isr}}{\sum_{s\in S,i\in I_{sr,su}}\left(\frac{\bar{v}_{isu}-\bar{v}_{isr}}{\ln \bar{v}_{isu}-\ln \bar{v}_{isr}}\right)} \tag{61}$$

$$= 0 \tag{62}$$

The usefulness of this result is that P_{ur} can be observed from the data, as can λ_u and λ_r. If we substitute this into the equation above we get,

$$\log\left(\frac{v_{isr}/\lambda_r}{v_{isu}/\lambda_u}\right) = (1-\sigma)\log\left(\frac{p_{isr}}{p_{isu}}/P_{ur}\right) \tag{63}$$

Figure 7: Price Indexes and Probability of a Sale

(a) Carb. Beverages — Price Indexes

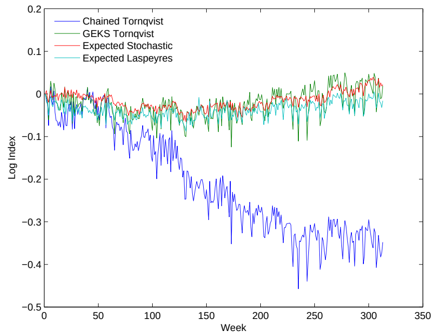

(b) Carb. Beverages — Probability of a Sale

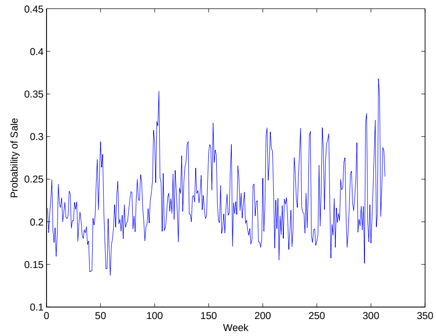

(c) Coffee — Price Indexes

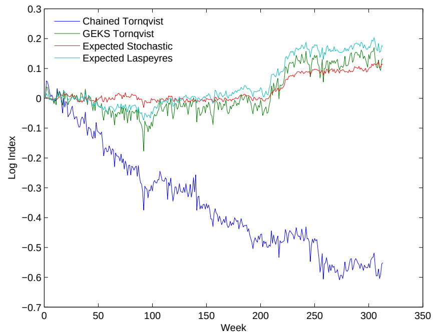

(d) Coffee — Probability of a Sale

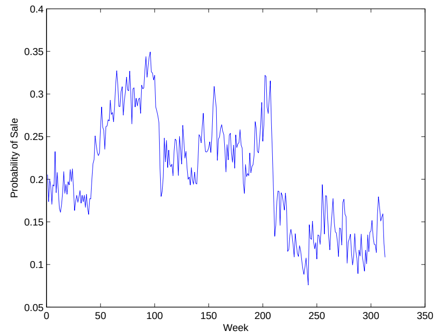

(e) Laundry Detergent — Price Indexes

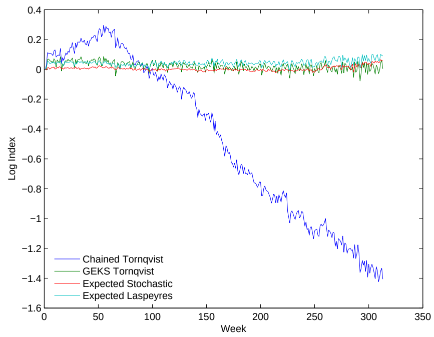

(f) Laundry Detergent — Probability of a Sale

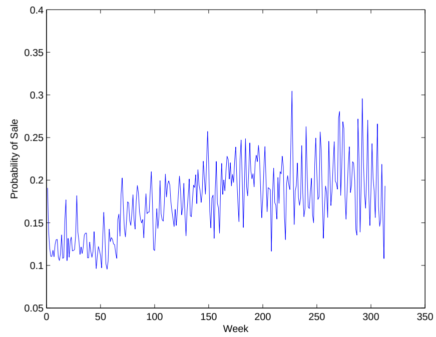

(g) Milk — Price Indexes

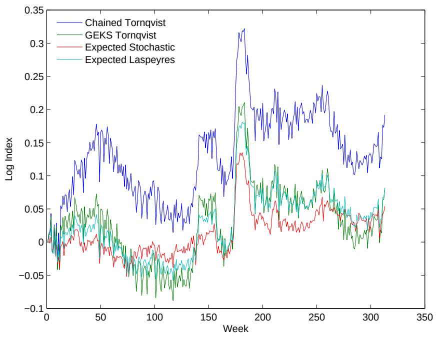

(h) Milk — Probability of a Sale

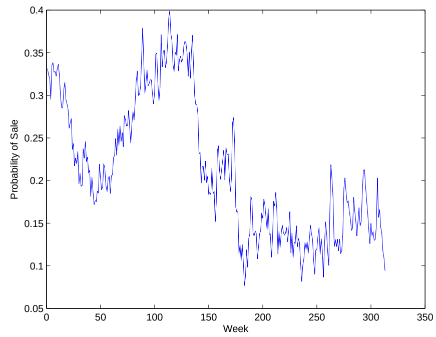

Figure 8: Price Indexes and Probability of a Sale

(a) Mustard and Ketc. — Price Indexes

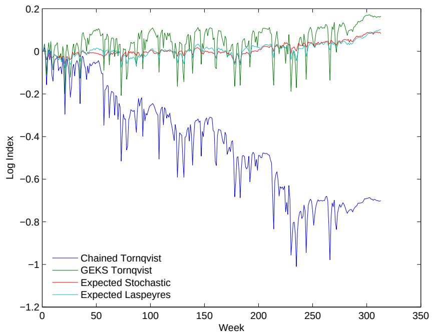

(b) Mustard and Ketc. — Probability of a Sale

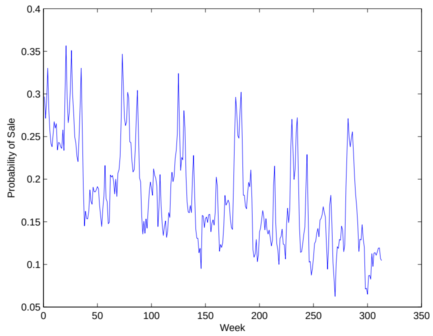

(c) Peanut Butter — Price Indexes

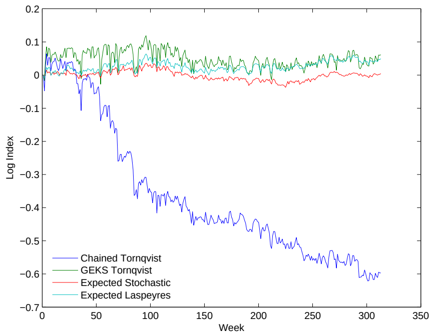

(d) Peanut Butter — Probability of a Sale

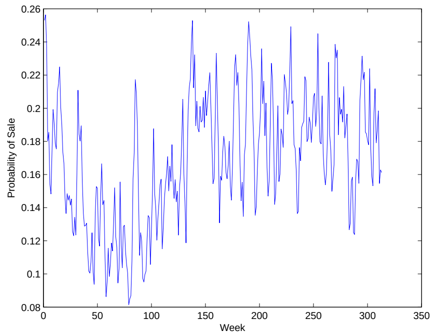

(e) Shampoo — Price Indexes

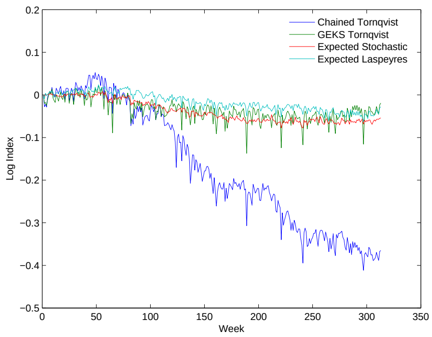

(f) Shampoo — Probability of a Sale

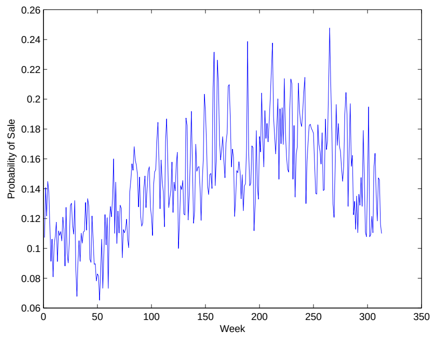

(g) Toothpaste — Price Indexes

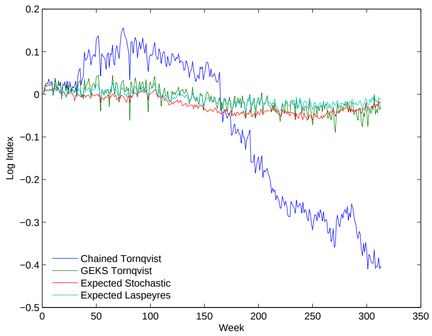

(h) Toothpaste — Probability of a Sale

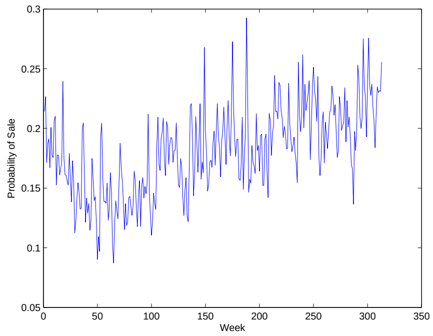

Figure 9: Multilateral Rolling Year Price Indexes

(a) Carbonated Beverages — Modeled Prices

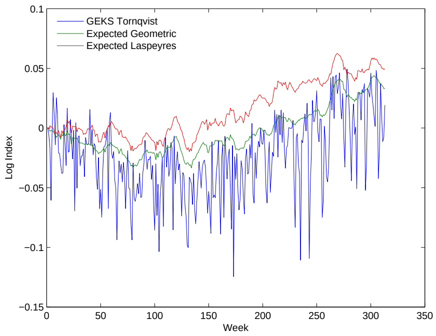

(b) Carbonated Beverages — Actual Prices

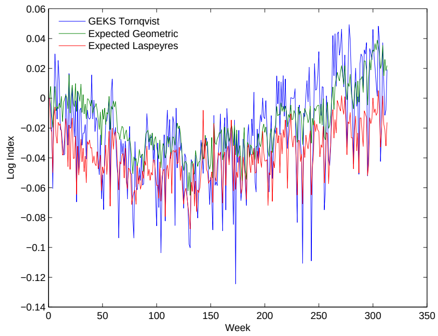

(c) Coffee — Modeled Prices

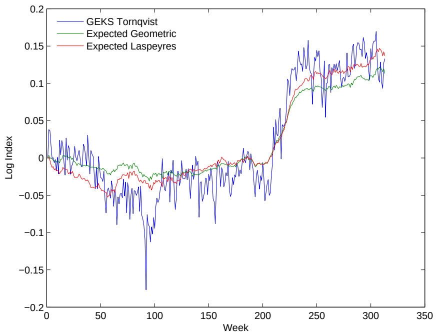

(d) Coffee — Actual Prices

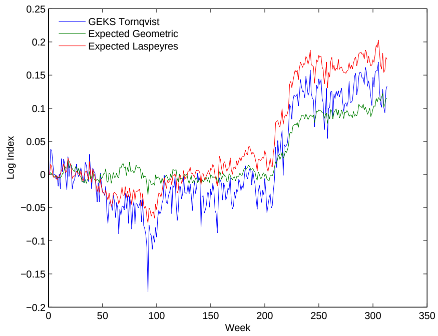

(e) Laundry Detergent — Modeled Prices

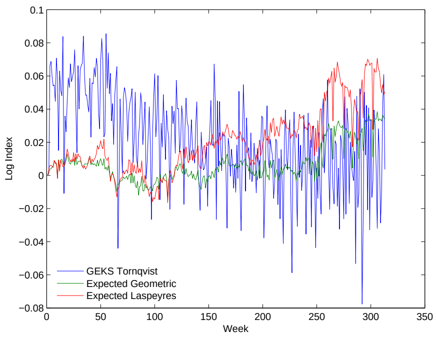

(f) Laundry Detergent — Actual Prices

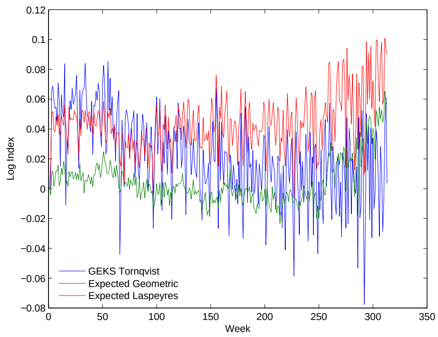

(g) Milk — Modeled Prices

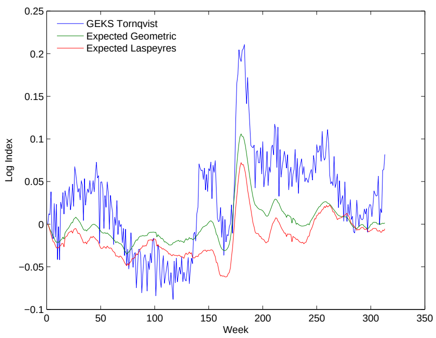

(h) Milk — Actual Prices

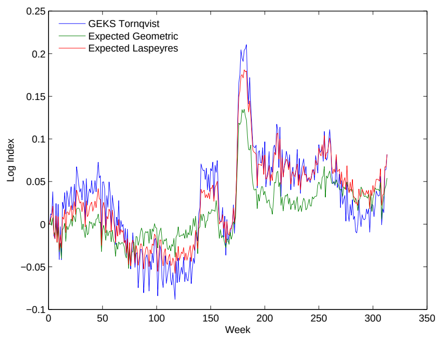

Figure 10: Multilateral Rolling Year Price Indexes

(a) Mustard and Ketchup — Modeled Prices

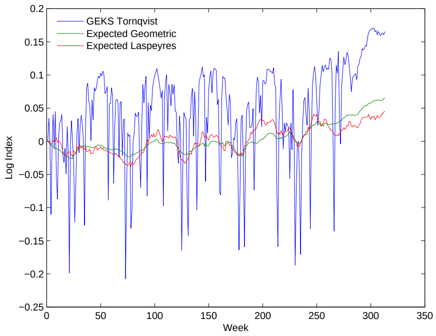

(b) Mustard and Ketchup — Actual Prices

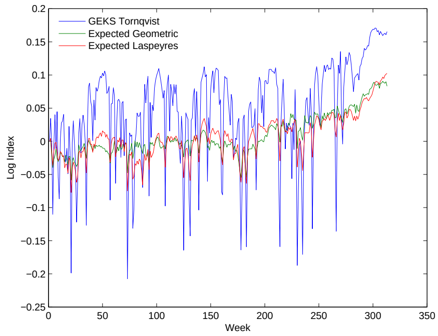

(c) Peanut Butter — Modeled Prices

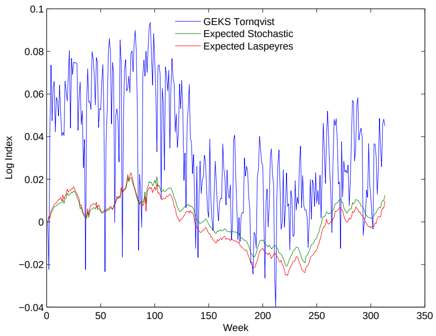

(d) Peanut Butter — Actual Prices

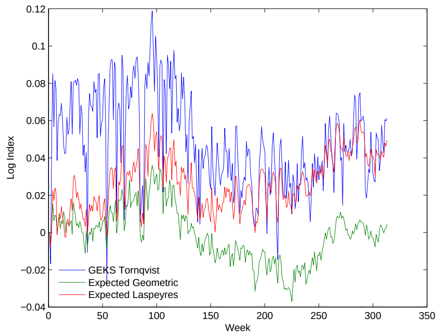

(e) Shampoo — Modeled Prices

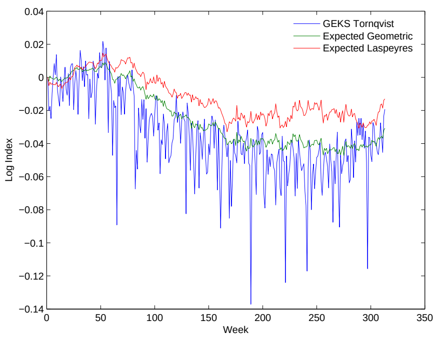

(f) Shampoo — Actual Prices

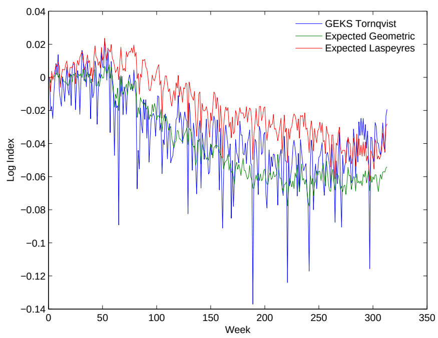

(g) Toothpaste — Modeled Prices

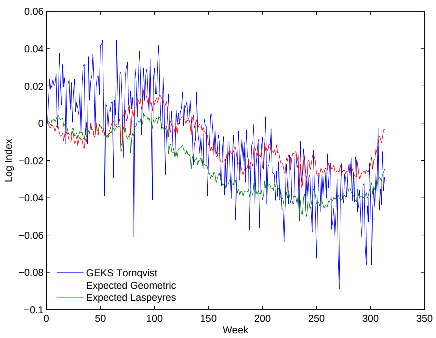

(h) Toothpaste — Actual Prices