

Fox, Kevin J.; Levell, Peter; O'Connell, Martin

Working Paper

Inflation measurement with high frequency data

IFS Working Papers, No. 23/29

Provided in Cooperation with:

Institute for Fiscal Studies (IFS), London

Suggested Citation: Fox, Kevin J.; Levell, Peter; O'Connell, Martin (2023) : Inflation measurement with high frequency data, IFS Working Papers, No. 23/29, Institute for Fiscal Studies (IFS), London, <https://doi.org/10.1920/wp.ifs.2023.2923>

This Version is available at:

<https://hdl.handle.net/10419/284255>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Kevin J. Fox
Peter Levell
Martin O'Connell

23/29

Working paper

Inflation measurement with high frequency data

Inflation measurement with high frequency data

Kevin J. Fox, Peter Levell, Martin O’Connell*

October, 2023

Abstract

The availability of large transaction level datasets, such as retail scanner data, provides a wealth of information on prices and quantities that national statistical institutes can use to produce more accurate, timely, measures of inflation. However, there is no universally agreed upon method for calculating price indexes with such high frequency data, reflecting a lack of systematic evidence on the performance of different approaches. We use a dataset that covers 178 product categories comprising all fast-moving consumer goods over 8 years to provide a systematic comparison of the leading bilateral and multilateral index number methods for computing month-to-month inflation.

Keywords: Consumer price index (CPI), chain drift, multilateral indexes, scanner data

JEL classification: C43, E31

Acknowledgements: The authors gratefully acknowledge financial support from the Office for National Statistics (ONS) as part of the research programme of the Economic Statistics Centre of Excellence (ESCoE) and the ESRC Centre for the Microeconomic Analysis of Public Policy at IFS (ES/T014334/1). We thank Erwin Diewert, Jan de Haan, Xavier Jaravel and Shipei Zeng, and the ONS Advisory Panel on Consumer Price Statistics for useful discussions. We make use of the R package IndexNumR by Graham White. His assistance is gratefully acknowledged. The contents of this publication are the authors’ own analysis and findings. The analysis and findings have been undertaken by authors’ on the basis of their own work using data supplied by Kantar FMCG At-Home Purchase Panel (all use is subject to Kantar UK Limited’s terms and conditions) and Kantar does not represent any endorsement of the efficacy or accuracy of the analysis and findings. The use of Kantar FMCG At-Home Purchase Panel data in this work does not imply the endorsement of Kantar FMCG Purchase Panel from the interpretation or analysis of the data. All errors and omissions remain the responsibility of the authors.

*Fox is at UNSW Sydney, CAER and ESCoE, Levell is at the Institute for Fiscal Studies and ESCoE and O’Connell is at the University of Wisconsin-Madison and Institute for Fiscal Studies.

Correspondence: k.fox@unsw.edu.au, peter.l@ifs.org.uk, moconnell9@wisc.edu.

1 Introduction

The ability to detect changes in the inflation rate in an accurate and timely way is essential for effective policymaking. For instance, central banks rely on inflation measures when setting interest rates, while their ability to maintain price stability depends on inflation expectations and inflation-indexed labor contracts, both of which are likely to be influenced by the most recent inflation statistics (e.g., Coibion et al. (2018)). In addition, benefit and social insurance programs are typically indexed to inflation measures, meaning their ability to provide protection for individuals against adverse inflation shocks depends on whether they quickly and accurately track sudden changes in consumer purchasing power.

Traditionally, National Statistical Institutes (NSIs) collect price data for their consumer price indexes (CPIs) using in-person collectors, which yield a relatively small sample of price quotes and no product level expenditure information. This can result in month-to-month changes in the index that are noisy. In addition, index methods traditionally used to construct CPIs, that weight product aggregates using historical expenditure weights, are ill-suited to the task of inflation measurement using high-frequency data as they can result in an index that becomes rapidly unrepresentative of spending patterns over time, leading to substantial substitution bias.¹

The increasing availability of large, comprehensive transaction-level scanner datasets, which contain close-to real-time expenditure and price information for thousands of products across millions of transactions, has created new opportunities for accurately measuring month-to-month price changes for important segments of the economy.² These datasets are being increasingly used to produce price indexes by NSIs.³ However, challenges facing any index using this type of data include the typically large degree of product entry and exit, and volatile movements in prices and quantities. These can lead to biases in the index (Ivancic et al. (2011), de Haan and van der Grient (2011)). Moreover, at present, there is no universally agreed upon method for calculating price indexes with high-frequency transaction data.

¹Substitution bias occurs when relative price changes cause consumers to change their consumption choices, yet expenditure weights representing, e.g., a previous period's expenditure patterns are used to weight the current period's prices.

²Scanner data sets typically cover fast-moving consumer goods, which make up approximately 40% of household expenditure on goods and 15% of expenditure on goods and services (see Jaravel (2019)).

³For instance, retail chains have agreed to share their high frequency product level volume and sales data with NSIs in several countries, including Australia, Canada, Japan, Netherlands, Norway, Switzerland and the UK (Diewert (2022)), with several of these agencies working on, or already, incorporating high-frequency transaction data into their CPI.

In this paper we provide a systematic empirical comparison of alternative index number methods for computing month-to-month inflation with high frequency transactions data. We use a scanner dataset from the Kantar FMCG At-Home Purchase Panel that covers fast-moving consumer goods, including food and drink (both alcoholic and non-alcoholic) products and household supplies, such as toiletries, non-prescription drugs, cleaning products, and pet foods. The dataset contains information on expenditures and transaction prices for over 290,000 unique products and over 300 million transactions over eight years (2012-19). We compute price indexes both for fast-moving consumer goods as a whole, as well as for each of the 178 product categories that together comprise this segment of the economy. We compare fixed base, chained bilateral and multilateral methods, in order to determine best practice.

We start by showing that the use of fixed weights can lead to a price index quickly becoming unrepresentative: only 32% of expenditure in the final month of our data is on products that were purchased in the first month. Chained bilateral indexes allow for the basket of products to be updated over time, but can introduce chain drift bias. This bias can arise in chained indexes when changes in product weights between two periods are correlated with prices in other periods (see Reinsdorf (1998) and Diewert (2022)). One reason this can happen is inventory stockpiling during sales, which can lead to product quantity and budget shares being lower immediately following the sale than prior to it. This bias results in the index failing a *multi-period identity test* – if all prices and weights return to their initial period values the chained price index will not equal its initial value.

We show that in practice chained bilateral indexes do exhibit substantial chain drift. For instance, a chained Törnqvist index, the type of index used by the US Bureau of Labor statistics in calculating the C-CPI-U measure of inflation, reports cumulative inflation of -16% for all fast-moving consumer goods over 2012-19, and cumulative inflation above 20% for 9 product categories, and below -50% for 10 product categories. Price changes of this magnitude for this sector are not realistic over this time period. For indexes computed at a higher level of time aggregation, chain drift bias is reduced but can remain significant; for instance, a quarter-to-quarter chained Törnqvist index for all fast-moving consumer goods records cumulative inflation of -7%, which is closer to, but remains below, the chain-drift-free multilateral index methods that we consider.

One approach to tackling this bias is to chain across periods that are not necessarily adjacent but are the most similar in terms of price structures or available products. This can in principle reduce chain drift bias, but the empirical perfor-

mance of such approaches has not been extensively explored, with multiple options available for deciding the order of chaining. We show that the bias associated with monthly chaining is mitigated only to a very limited extent when using this approach. While in principle this approach should reduce chain drift bias, we find that in practice the immediately preceding months are often selected as the most similar, and as a result this approach yields similar results to chaining period-on-period.

A second approach to tackling chain drift bias is to use multilateral index methods, which compare the current month’s price level to all previous months over which the index is computed and which have no chain drift bias. We compare three multilateral indexes that take the Törnqvist, Fisher and Walsh indexes as their basic building blocks. These bilateral indexes are all “superlative”; a superlative index formula is exactly equal to a Könus (1924) true cost of living index if households have preferences that can be represented by certain functional forms, where these functional forms can approximate arbitrary preferences to the accuracy of a second-order approximation (Diewert (1976)). The corresponding multilateral indexes are called Caves-Christensen-Diewert-Inklaar (CCDI), GEKS-Fisher and GEK-Walsh indexes (where GEKS stands for Gini-Eltető-Köves-Szulc).

Using our dataset, the CCDI index reports cumulative inflation for all fast moving consumer goods of 2.5%. Both for the segment as a whole and individual product categories, the GEKS indexes generate similar results to the CCDI index, though we find the former are more sensitive to large product level price changes likely to reflect measurement error. We also consider the Geary-Khamis (GK) index. Unlike the other multilateral indexes we consider, the GK index is not built from a superlative bilateral index, but instead is microfounded by a consumer model with either linear or Leontief preferences, making the index more susceptible to substitution bias. We find that this is borne out in practice; across product categories the 25th and 75th percentiles of the distribution of differences in average monthly inflation with the CCDI index are -0.002 percentage points (ppt) and 0.002 for the GEKS-Fisher, and -0.004 ppt and 0.004 ppt for the GEKS-Walsh indexes, but -0.01 ppt and 0.02 ppt for the GK index.⁴

A practical drawback of multilateral indexes, in their pure form, is they entail revisions to historic numbers as each new month of data becomes available, mak-

⁴Another approach to addressing chain drift is to compute a sequence of year over year monthly indexes. This approach was used by Handbury et al. (2013), and has the benefit of avoiding seasonality problems. Our focus is on methods for calculating month-to-month CPIs, consistent with NSI practice and the needs of central banks and other users.

ing them inappropriate for use in CPIs.⁵ Splicing methods, which link multilateral indexes calculated in different ‘windows’ of time, eliminate the need to make CPI revisions, but reintroduce some chain drift bias. While there have been various suggestions by researchers and practitioners, and some NSIs have implemented different options, the performance of different splicing methods is still debated and is ultimately an empirical question.

We evaluate alternative splicing methods and window lengths by comparing spliced series to their corresponding chain-drift-free multilateral index series computed over all 96 months of our data; the difference provides a quantification of the chain drift bias the splicing procedure induces. We show that spliced multilateral indexes tend to exhibit much smaller chain drift bias than their chained bilateral counterparts. Nonetheless, chain drift bias in spliced series can be significant and therefore the choice of index number formula and splicing procedure matter in practice. We find the differences in chain drift biases across splicing methods for the CCDI and GEKS indexes are not large, but the degree of chain drift bias exhibited by the Geary-Khamis index is highly sensitive to the choice of splicing method.

Our results suggest the CCDI multilateral index as the preferred option, using a 25 month window length. While the choice of splicing method appears relatively unimportant for the degree of chain drift bias, we argue for the use of the mean splice which averages over different splicing periods and so is less sensitive to the risk of linking on a particularly anomalous month.

We conclude our analysis by exploring what the main drivers of chain drift bias for spliced multilateral indexes are. We show that the most significant predictor of chain drift bias in spliced series across product categories is the degree of product entry and exit, or “churn”, in the category. Monthly and annual product churn contribute to 88% of the total explained variance in chain drift in 25 month spliced CCDI indexes, with the rest being due to the frequency of price and quantity promotions and seasonality in prices. This analysis illustrates that categories with rapid product turnover are where chain drift bias is likely to be most severe, and thus where longer window lengths are likely to be of most benefit.

Our work adds to a literature that evaluates different index number methods for measuring inflation with high frequency scanner data. This includes Ivancic et al. (2011) who use scanner data on a set of 19 product categories over 15 months to quantify the extent of chain drift among superlative bilateral indexes, through comparison with a chain drift free GEKS-Fisher index and Diewert and Fox (2022)

⁵For instance, the US Bureau of Labor Statics and the UK NSI, the Office for National Statistics, have a policy of never revising published headline CPI numbers, unless a significant error has occurred. This is standard practice internationally.

who simulate a dataset based on consumers with constant elasticity of substitution (CES) preferences and compare price series computed with different bilateral and multilateral index numbers to the true CES cost-of-living index.

A limitation of previous work is that empirical evidence has been limited to a subset of available index methods and a small number of product categories over relatively short time periods. We contribute to this literature by systematically comparing the set of leading methods for measuring inflation using transaction data of significantly broader scale and scope.

Our work also contributes to the broader literature that uses scanner data to advance understanding of various aspects of inflation measurement, including work quantifying the impact of product entry and exit on changes in the cost-of-living (e.g., Broda and Weinstein (2010)), documenting a divergence in inflation experienced by consumers and in posted prices (e.g., Coibion et al. (2015)) and measuring the degree of heterogeneity in inflation rates across consumers (e.g., Kaplan and Schulhofer-Wohl (2017), Jaravel (2019)). We contribute to the strand of this literature focused on the measurement of high frequency price dynamics, and that includes work documenting the frequency of price adjustment of individual products (e.g., Eichenbaum et al. (2011); see also the survey by Nakamura and Steinsson (2013)) and high frequency inflation during the COVID-19 pandemic (Jaravel and O’Connell (2020)).

The remainder of this article is structured as follows. In Section 2 we outline the different index number methods and in Section 3 we describe the dataset we use to empirically assess their performance. In Section 4 we compare fixed-based and chained bilateral indexes with multilateral indexes. In Section 5 we focus on spliced multilateral indexes, quantifying how their degree of chain drift bias varies with the linking method and window length used in their construction, and exploring the drivers of this bias. We conclude and discuss potential avenues for future research in a final section.

2 Inflation measurement with high-frequency data

Suppose, for a sequence of periods $1, \dots, T$, we observe period specific prices, $\mathbf{p}^t = (p_1^t, \dots, p_N^t)'$, and quantities $\mathbf{q}^t = (q_1^t, \dots, q_N^t)'$ for N goods, and we wish to compare how the cost of purchasing the basket of goods changes over time. In this section we examine alternative ways of doing this, with particular reference to the use of high-frequency (scanner/transactions) data.

2.1 Bilateral index numbers

Suppose we are interested in comparison of the change in the cost of the basket of goods between two sequential periods, t and $t+1$. One way of measuring this change is with a Lowe price index, which takes the form $P_{Lo}^{t,t+1} = \frac{\mathbf{p}^{t+1'} \mathbf{q}}{\mathbf{p}^t \mathbf{q}}$, and is commonly used in CPI construction. If base period quantities are used (i.e., $\mathbf{q} = \mathbf{q}^t$) the index is known as a Laspeyres index ($P_L^{t,t+1}$) and if end period quantities are used ($\mathbf{q} = \mathbf{q}^{t+1}$) it is known as Paasche index ($P_P^{t,t+1}$). The indexes can be re-written in terms of price relatives for each good between t and $t+1$, $\frac{p_n^{t+1}}{p_n^t}$, weighted by the share of expenditure allocated to them, $s_n^t = \frac{p_n^t q_n^t}{\mathbf{p}^t \mathbf{q}^t}$:

$$P_L^{t,t+1} = \sum_n s_n^t \frac{p_n^{t+1}}{p_n^t}$$

$$P_P^{t,t+1} = \left(\sum_n s_n^{t+1} \left(\frac{p_n^{t+1}}{p_n^t} \right)^{-1} \right)^{-1}.$$

A drawback of these indexes is that they suffer from substitution bias. This bias arises as the indexes use weights corresponding to just one (either the base or end) period, and therefore fail to reflect that when the price of goods increase/decrease consumer typically substitute away/towards them. This leads to upward bias in the case of the Laspeyres index and downwards bias in the case of the Paasche index.

A solution to the problem of substitution bias is offered by *superlative* indexes (see Diewert (1976)). These use a combination of base and final period weights. Three commonly used superlative indexes are the Fisher index (a geometric mean of the Laspeyres and Paasche indexes), $P_F^{t,t+1}$, the Törnqvist index (a geometric mean of price changes weighted by average spending shares in the base and end periods), $P_{Tq}^{t,t+1}$, and the Walsh index (an arithmetic average of price changes weighted by the geometric mean of quantities in the base and end periods), $P_W^{t,t+1}$:

$$P_F^{t,t+1} = (P_L^{t,t+1} P_P^{t,t+1})^{1/2}$$

$$P_{Tq}^{t,t+1} = \prod_n \left(\frac{p_n^{t+1}}{p_n^t} \right)^{0.5(s_n^t + s_n^{t+1})}$$

$$P_W^{t,t+1} = \frac{\sum_n \sqrt{q_n^t q_n^{t+1}} p_n^{t+1}}{\sum_n \sqrt{q_n^t q_n^{t+1}} p_n^t}$$

2.2 Chaining and chain drift

Now suppose we are interested in how the cost of the basket of goods changes across several periods, and consider the comparison of the first and some subsequent

period $s > 2$. One way of making this comparison is by constructing the chain of intervening period-to-period comparisons, $P^{1,2} \times P^{2,3} \times \dots \times P^{s-1,s}$, which results in a *chained* index. An alternative is to make the direct, or *fixed-base*, comparison $P^{1,s}$. Notice that in this case the comparison of prices in two sequential periods is given by $P^{1,t+1}/P^{1,t}$. A major limitation of fixed-based indexes arises in contexts in which there is product churn (entry and exit of products through time), as in this case the comparison of price levels across two periods is between continuing products (i.e., those available in both periods), and in the case of the fixed-base index the set of overlapping products can fall rapidly over time.⁶ Hence, there is an international consensus that it is preferable to use chained indexes as they more closely follow market-place developments.⁷

A drawback of chained indexes is that (unlike fixed-base indexes) they can exhibit *chain drift*. Suppose that the period T price and quantity vectors coincide exactly with the price and quantity vectors in period 1 (i.e., by the end of the period price and quantities return to their initial level). A chained index is said to satisfy the multiperiod identity test if $P^{1,2} \times P^{2,3} \times \dots \times P^{T-1,1} = 1$; i.e., if the chained comparisons between the first and final period – identical in terms of prices and quantities – leads to the conclusion that the two periods have the same price level. In general, bilateral chained indexes (including those constructed with the five index numbers defined above) fail this test and are therefore said to suffer from chain drift. The resulting bias when indexes are computed with transaction level data, like scanner data, can be severe (e.g., Ivancic et al. (2011)).

Diewert (2022) shows that chain drift bias arises with a Törnqvist index if the change in spending shares between two periods are correlated with prices in any other period. This will be the case, for instance, if, when a product goes on sale, consumers stock up so that when price returns to the regular level spending shares do not, as consumers draw down their stocks. From the economic approach to index numbers (Diewert (1976)) the Törnqvist index is derived from second order approximations to a consumer expenditure function defined over consumption (with the same being true for Fisher and Walsh indexes). Yet, when consumers stockpile, consumption does not align with transactions, the information used to compute

⁶Consider, comparison of period t and $t + 1$. The chained comparison $P^{t,t+1}$ will use data on the set of products available in periods t and $t + 1$. The fixed-base comparison, $P^{1,t+1}/P^{1,t}$, will use data on the set of products available in both periods 1 and t or 1 and $t + 1$. If there is entry and exit, the set of products available in all periods 1 and t (or $t + 1$) may comprise a much smaller share of spending than those available in periods t and $t + 1$, especially when t is several periods after the initial period.

⁷For instance, the ILO (2004) CPI Manual (p. 407) states: “rapid sample attrition means that fixed-base indexes rapidly become unrepresentative and hence it seems preferable to use chained indexes which can more closely follow market-place developments.”

the price index. Chain drift bias can also arise when there is product churn. For example, suppose a product's price either rises or falls (due to reduced availability or falling popularity), but its spending share is stable, in the period prior to its exit from the market. A fixed based comparison between the post-exit and an initial period (based on overlapping products) will be unaffected by the exiting product's price trend, whereas this is not true for a chained comparison.

Dissimilarity chain linking. One approach to minimizing chain drift bias, which uses bilateral index number comparisons, is to construct chains of bilateral indexes across periods that are closest to having proportional increases or decreases in prices from each other. This requires a measure of ‘dissimilarity’ of price levels between periods.⁸ Diewert et al. (2022) recommends a ‘predicted share measure of relative price dissimilarity’ for the calculation of price indexes in cases where there is a high rate of product churn.⁹ This approach is motivated by the *multi-period* identity test: bilateral indexes will satisfy this property when price changes across periods are strictly proportional.

This measure uses hypothetical expenditure shares using prices in period τ but quantities from period t to ‘predict’ expenditure shares in period t . For any good n , this predicted share is

$$\tilde{s}_{n,t,\tau} = \frac{p_n^\tau q_n^t}{\mathbf{p}^{\tau'} \mathbf{q}^t}$$

The predicted share measure of relative price dissimilarity between periods t and τ is then

$$\Delta_{PS}(p^t, p^\tau, q^t, q^\tau) \equiv \sum_{n=1}^N [s_{n,t} - \tilde{s}_{n,t,\tau}]^2 + \sum_{n=1}^N [s_{n,\tau} - \tilde{s}_{n,\tau,t}]^2$$

This measure takes on values between 0 and 2. It takes the value 0 if prices in period τ are proportional to prices in period t (i.e. $\mathbf{p}^\tau = \lambda \mathbf{p}^t$), as in this case $s_{n,t} = \tilde{s}_{n,t,\tau}$ and $s_{n,\tau} = \tilde{s}_{n,\tau,t}$ for all n .

This approach seeks to limit the degree of chain drift bias by only linking across those periods that are closest to having proportional increases or decreases in prices from each other. However, as we find in Section 4, it can still be associated with substantial biases.

⁸This approach is also used in the context of international comparisons, identifying country pairs that are most similar to each other. See Hill (1999) and Hill (2001).

⁹Diewert (2002) sets out a number of alternative possible dissimilarity measures.

2.3 Multilateral index numbers

Another solution to chain drift bias is offered by multilateral index numbers, which were first suggested as a solution to chain drift bias by Ivancic et al. (2011).¹⁰ A multilateral index computed over all periods $1, \dots, T$ will satisfy the multiperiod identify test, and does not suffer from the problem of the basket of products becoming increasingly unrepresentative over time as much as fixed-base indexes. We consider four different multilateral indexes. The first three use the superlative bilateral indexes defined in Section 2.1 (i.e., the Fisher, Törnqvist and Walsh indexes) as their building blocks, and thus like them are consistent with a flexible representation of consumer preferences and hence limit substitution bias. They are called the GEKS-Fisher index, CCDI index and the GEKS-Walsh index respectively.¹¹ In each case, the price level in period t is given by a geometric mean of the corresponding bilateral index that compares period t with all other periods $\tau = 1, \dots, T$. Hence, the measured price level in period t under the indexes is given by:

$$\begin{aligned}\mathbb{P}_{GEKS-F}^t &= \prod_{\tau} [P_F^{\tau,t}]^{1/T} \\ \mathbb{P}_{CCDI}^t &= \prod_{\tau} [P_{Tq}^{\tau,t}]^{1/T} \\ \mathbb{P}_{GEKS-W}^t &= \prod_{\tau} [P_W^{\tau,t}]^{1/T}.\end{aligned}$$

The fourth multilateral index number we consider is the GK index,¹² which is an implicit price index, defined as total expenditure divided by a volume or quantity index, with ‘quality adjustment factors’ determining how many units of good m are equivalent to a unit of good n . Unlike the other multilateral indexes we consider, the GK index is not based on an underlying superlative bilateral index. Rather, it is based on a linear preference model in which consumers view goods as perfect substitutes. Diewert and Fox (2022) show the index is also consistent with Leontief – or perfect complement – preferences. That is, it is only consistent with extreme assumptions on consumer behaviour. The index is implicitly defined by the solution to a set of equations that jointly determine price levels, \mathbb{P}_{GK}^t , for $t = 1, \dots, T$ and quality adjustment factors, b^n for $n = 1, \dots, N$. It is helpful to denote the total

¹⁰Multilateral indexes are typically used in contexts such as international comparisons, such as the World Bank’s International Comparisons Program (<https://www.worldbank.org/en/programs/icp>). They were first suggested in a time series context by Balk (1981).

¹¹The GEKS indexes are named after Gini (1931), Eltetö and Köves (1964) and Szulc (1964) and the Caves-Christensen-Diewert-Inklaar (CCDI) index was developed by Caves et al. (1982) and applied to a price index by Inklaar and Diewert (2016).

¹²Developed in Geary (1958) and Khamis (1970, 1972).

quantity of good n across all time periods by $q_n \equiv \sum_t q_n^t$. The $N + T$ equations that determine the quality adjustment factors and price levels are:

$$b_n = \sum_t \left(\frac{q_n^t}{q_n} \right) \left(\frac{p_n^t}{\mathbb{P}_{GK}^t} \right) \quad \text{for } n = 1, \dots, N$$

$$\mathbb{P}_{GK}^t = \frac{\mathbf{p}'^t \mathbf{q}^t}{\mathbf{b}' \mathbf{q}^t} \quad \text{for } t = 1, \dots, T.$$

Each adjustment factor b_n is a share-weighted average of inflation-adjusted prices for each commodity n over all t periods.¹³ The final index is expenditure divided by the sum of quality adjusted quantities purchased in each period.

For each of the multilateral index numbers, it is common to rebase the price levels relative to the first period of date: $P_i^{1,t} = \mathbb{P}_i^t / \mathbb{P}_i^1$, for $i = GEKS-F, CCDI, GEKS-W, GK$. The comparison of prices in period t and $t+1$ is given by $P_i^{t,t+1} = \mathbb{P}_i^{t+1} / \mathbb{P}_i^t = P_i^{1,t+1} / P_i^{1,t}$.

2.4 Spliced price series

Suppose we use a multilateral index to compute price levels over a given time period, $1, \dots, T$. If data for period $T + 1$ becomes available, re-computing the index over $1, \dots, T + 1$ will lead to a revision of price levels over the initial T periods. NSIs regard such revisions to past headline CPI levels as undesirable. Linking methods provide a way to avoid this problem. A number of alternative methods for splicing multilateral indexes calculated over rolling windows have been proposed, as follows.

The idea behind the rolling window splice is to compute an initial multilateral index over periods $t = 1, \dots, T$. When a new period of data becomes available, a new sequence can be computed over periods $t = 2, \dots, T + 1$. The price level for period $T + 1$ computed with this new sequence can then be spliced to the original series using the price levels computed in the two series for some chosen comparison period. As each new period of data arrives, a new sequence of length T is computed and this is used to splice the new data point to the spliced sequence. In this example, the window length is given by T .

More concretely, suppose we have a multilateral price series computed over $t = 1, \dots, T$, $\mathbb{P}_{\mathcal{O}} = (\mathbb{P}_{\mathcal{O}}^1, \dots, \mathbb{P}_{\mathcal{O}}^T)$. For $t \leq T$, the price level is $\rho_t = \frac{\mathbb{P}_{\mathcal{O}}^t}{\mathbb{P}_{\mathcal{O}}^1}$. In the next period, $T + 1$, we then compute a new multilateral sequence over the periods $t = 2, \dots, T + 1$, $\mathbb{P}_{\mathcal{N}} = (\mathbb{P}_{\mathcal{N}}^2, \dots, \mathbb{P}_{\mathcal{N}}^{T+1})$. The spliced price level for period $T + 1$ is given

¹³The usual method for obtaining a solution to these two equations is to iterate between them. However, Diewert and Fox (2022) derive an alternative method which is more efficient (p. 360, footnote 24).

by:

$$\rho_{T+1}(\tau) = \rho_T(\tau) \times \frac{\mathbb{P}_{\mathcal{N}}^{T+1}/\mathbb{P}_{\mathcal{N}}^{\tau}}{\mathbb{P}_{\mathcal{O}}^T/\mathbb{P}_{\mathcal{O}}^{\tau}},$$

where τ is the period used to link the series together. Different choices of τ correspond to different forms of the rolling-window splice;

- $\tau = T$ is known as the *movement splice* (Ivancic et al. (2011))
- $\tau = 2$ is known as the *window splice* (Krsinich (2016))
- $\tau = \frac{T}{2}$ (or, when T is an odd number $\tau = \frac{T+1}{2}$) is known as the *half splice* (de Haan (2015))

As each subsequent period of data, $t = s + T$ (for $s > 0$), arrives, the most recent T length multilateral sequence, $\mathbb{P}_{\mathcal{N}'} = (\mathbb{P}_{\mathcal{N}'}^{t-T+1}, \dots, \mathbb{P}_{\mathcal{N}'}^t)$ is added to the spliced series via the preceding period T length sequence $\mathbb{P}_{\mathcal{O}'} = (\mathbb{P}_{\mathcal{O}'}^{t-T}, \dots, \mathbb{P}_{\mathcal{O}'}^{t-1})$ and spliced price level $\rho_{t-1}(\tau)$. Hence,

$$\rho_{t+1}(\tau) = \rho_t(\tau) \times \frac{\mathbb{P}_{\mathcal{N}'}^{t+1}/\mathbb{P}_{\mathcal{N}'}^{\tau+s}}{\mathbb{P}_{\mathcal{O}'}^t/\mathbb{P}_{\mathcal{O}'}^{\tau+s}},$$

Without structure on the underlying price and quantity data, there is no obvious reason for favoring any $\tau = 2, \dots, T$. Rather than selecting one period, the *mean splice* (Diewert and Fox (2022)) uses a geometric mean of all possibilities, leading to normalized price level for calendar time $t > T$:¹⁴

$$\rho_t(\bar{\tau}) = \prod_{\tau=2}^T (\rho_t(\tau))^{\frac{1}{T-1}}$$

A final option considered is to select the splicing period using a *dissimilarity* measure such as the predicted share measure of relative price dissimilarity discussed above (i.e., setting $\tau = \arg \min_{\tau \in 2 \dots T} \Delta_{PS}(p^T + 1, p^{\tau}, q^T + 1, q^{\tau})$). This identifies the splicing period that is closest to being a proportional price change from the final period of the new window.¹⁵ This type of approach was suggested, but not pursued, by Diewert and Fox (2022). Hence, this paper presents the first evidence of its empirical performance.

¹⁴The idea of using a mean splice was originally suggested, but not pursued, by Ivancic et al. (2011), footnote 19, p. 33.

¹⁵An alternative would be to select the splicing period that is most similar to the final period of the old window, i.e., period T . This is likely to yield very similar results when constructing monthly indexes in practice however, as in most cases months T and $T + 1$ will have similar price structures.

Alternative splicing methods include the fixed-base moving and expanding windows and methods that splice on the published series (Chessa (2021)). We discuss these approaches and their empirical performance in Appendix A.

All of these linking procedures avoid the need to revise past price levels. However, this comes at the cost of introducing chain drift into the price index. The extent of the resulting bias depends on the window length and linking method chosen, and the nature of the underlying price and quantity data. While *a priori* it seems likely that a shorter window length will result in more bias, it is an empirical question to what extent this is true. In addition, without empirical evidence, it is unclear which linking method will perform best in practice. We provide this evidence by comparing a multilateral index computed over all periods in our data with the same index number computed with each of the linking procedures. The former satisfies the multiperiod identity test and therefore does not suffer from chain drift bias. The difference between the series provides an empirical quantification of the extent to which linking introduces chain drift bias into the price index.

3 Scanner data

We use household scanner data from the Kantar FMCG At-Home Purchase Panel. The data cover purchases of all fast-moving consumer goods (FMCG) – food, drink (including alcohol), toiletries, non-prescription drugs, cleaning products, and pet foods – brought into the home by a sample of households living in Great Britain (i.e., the UK excluding Northern Ireland). Our data cover the period 2012–2019. In each year, the dataset contains purchase records of around 30,000 households. Participating households are typically present in the data for many months. Each household records all barcodes that they purchase using a handheld scanner or mobile phone app. For each transaction we observe quantity, expenditure, transaction price and barcode characteristics (including product category).

In total, our data include 296,829 unique barcodes and over 300 million transaction, which are divided into product categories. In subsequent sections, we compute price indexes for each of the 178 product categories that account for at least 0.1% of total spending over 2012–2019.¹⁶ We compute monthly price indexes and treat barcodes as the elementary products in the index. We compute elementary (monthly

¹⁶We list the product categories in the Online Appendix. For each product category-year we drop transactions for which expenditure, volume or their ratio is in the top or bottom percentile; this does not affect our findings.

barcode) prices by dividing the total monthly expenditure for the barcode by the total monthly quantity.¹⁷

Our data are *household* scanner data, meaning they cover transactions made by a sample of households. In contrast, *retail* scanner (or point-of-sale) data cover all sales recorded in a sample of stores. Both data sources allow for close to real-time measurement of prices and expenditures for a comprehensive set of products in the fast-moving consumer good segment of the economy. Each type has its own advantages – retail scanner data is likely to have less sampling variation if store coverage is extensive, whereas household scanner data can include online purchases and allow for calculation of inflation for different types of households. In practice, there may be returns from combining both sources for inflation measurement.

Traditionally, NSIs obtain data on prices from quotes gathered in-person by collectors and combine this with survey data on expenditures (for instance, from the Consumer Expenditure Survey in the US and the Living Cost and Food Survey in the UK). These data have several limitations compared with scanner data, which are reflected in the standard approach to CPI construction. First, price quotes are collected for a limited set of products. Second, expenditure information is not available for disaggregate products, but at a level similar to the product categories in our scanner data. Third, the expenditure information is only available with a significant lag (typically at least one year). Therefore, CPIs traditionally entail taking unweighted averages over a relatively small number of price quotes to obtain product category level prices, which in turn are weighted based on historic spending patterns. Using large scanner datasets has many advantages in comparison, if the problem of chain drift can be mitigated.¹⁸

4 Comparing bilateral and multilateral indexes

One approach to calculating month-to-month indexes with high-frequency data is to use a fixed-base Laspeyres index. In common with traditional CPIs this index uses historic spending weights.¹⁹ However, when computed with scanner data, the

¹⁷That is, we use monthly unit values as the prices for index construction. Diewert et al. (2016) show that unit values should be calculated at the same frequency as the desired price index series to avoid introducing upward bias.

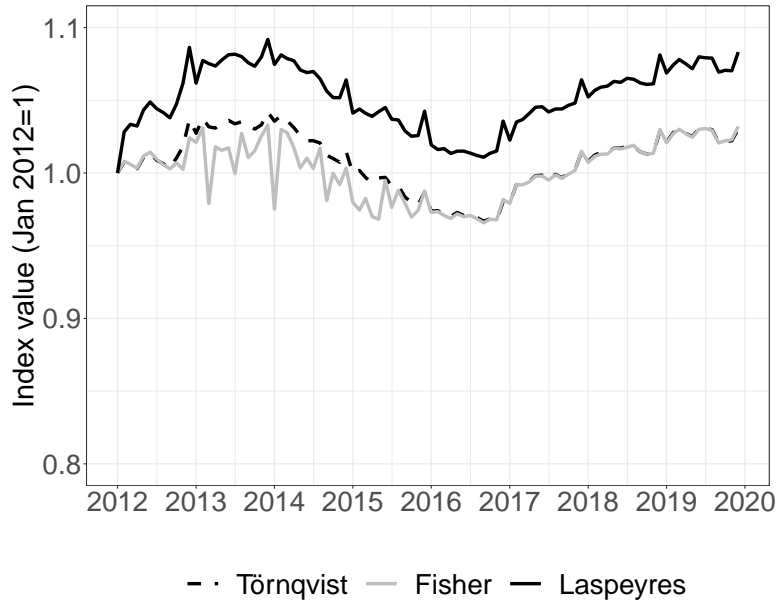
¹⁸In explaining their adoption of multilateral methods in constructing the CPI, the Australian Bureau of Statistics (2017) notes the following: “The advent of readily available transaction level data then allows for an overhaul of traditional methodology, as the data constraint has been enormously relaxed. However, this opportunity for improved price index construction has been somewhat offset by the complexities involved in the use of high-frequency data.”

¹⁹Most NSIs use a Laspeyres-type index in calculating the CPI, specifically the Lowe (1823) index. For CPI construction, the quantity weights in this index are typically from a period before

index has the advantage over traditional methods of being computed over thousands of products and including product level weights. Figure 4.1 plots the evolution of a fixed-base month-to-month Laspeyres index for all fast-moving consumer goods, computed over the period 2012-19. The figure also plots the fixed-base superlative Törnqvist and Fisher indexes – we omit the Walsh index, as the differences with the Törnqvist are very small.

The superlative indexes register substantially different price changes to the Laspeyres, which is indicative of substitution bias in the latter index. At the end of the first year of data there is a difference of around five percentage points between the Laspeyres and the superlative indexes, and this difference persists until the end of the sample period. In the first half of data, the Fisher index exhibits much more volatility than the Törnqvist and Walsh indexes, which is a theme of our results (including the multilateral extensions of these indexes) and reflects the sensitivity of the Fisher index to outlier price observations.

Figure 4.1: *Fixed-base Laspeyres and superlative indexes*



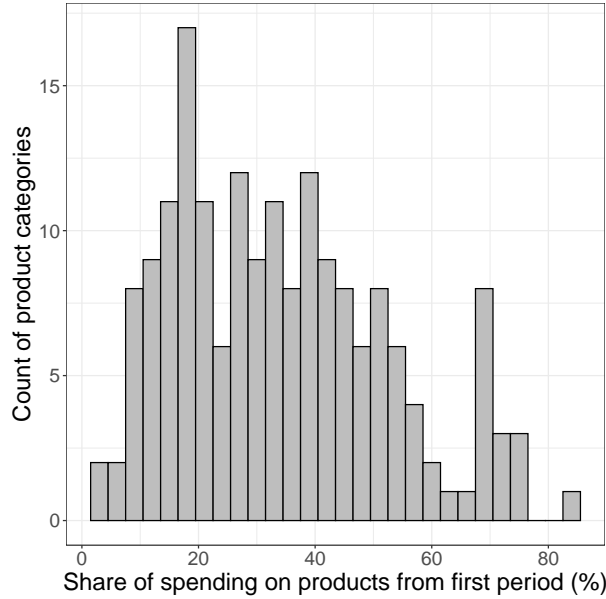
Note: Figure shows index number values for the Laspeyres, Törnqvist and Fisher fixed base indexes. We omit the Walsh index as it is very similar to the Törnqvist index. Indexes are computed over all fast-moving consumer goods.

All fixed-base indexes of this kind, whether superlative or not, run the risk of becoming increasingly unrepresentative of consumer spending as the availability of products changes across seasons and over time. Figure 4.2 is a histogram showing the distribution of the share of spending across the 178 product categories in Decem-

either of the periods of the price vectors being compared, usually the period of the last expenditure survey, or an imputed update of these weights.

ber 2019 allocated to products for which we observe positive spending in January 2012. The figure suggests there is significant product churn. For the median product category, 32% of spending in December 2019 went on items that had positive spending associated with them in January 2012.²⁰ Product churn on this measure is highest for moist wipes, machine wash products, cat food and fresh bacon joints.

Figure 4.2: *Share of final period spending on products available in first period*



Note: Figure shows the distribution of the share of spending in the final period (December 2019) that goes on products that were purchased in the first period (January 2012) across product categories.

Chained indexes help deal with product churn by only requiring products to be available in the two periods being compared in each bilateral link of the chain, and they also ensure that index weights reflect up-to-date spending patterns. However, as discussed in Section 2.2, chaining of bilateral indexes when using high frequency data can lead to significant chain drift. Figure 4.3(a) shows index values for three indexes, a fixed-base Törnqvist, a Törnqvist that is chained from month-to-month (‘period on period’) and the CCDI index, the multilateral analogue of the Törnqvist index, defined in Section 2.3. We calculate the indexes using all fast-moving consumer goods. We include equivalent graphs for the Fisher and Walsh indexes, which exhibit similar patterns, in the Online Appendix.

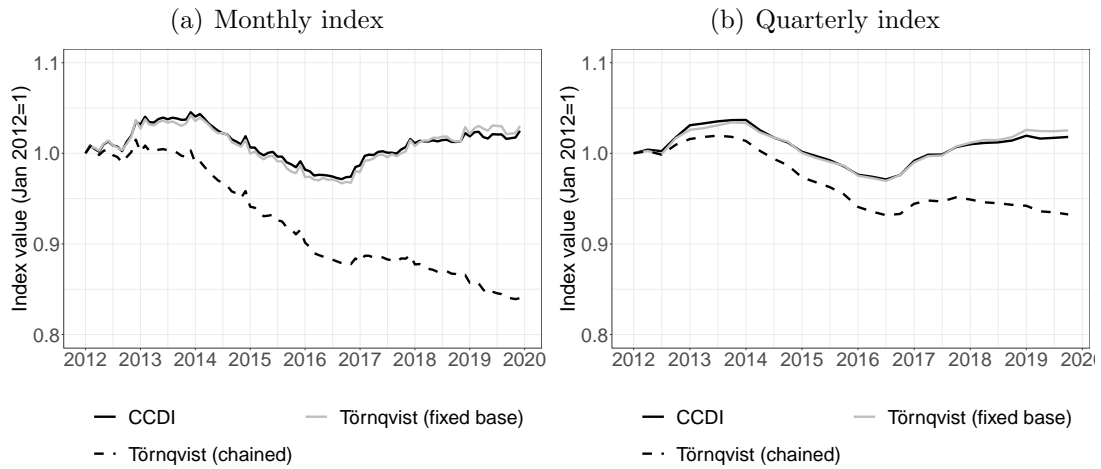
Figure 4.3(a) shows that chain drift is a significant problem for the Törnqvist index. The chained Törnqvist index is 19 percentage points lower at the end of the period than the fixed-base index and 18 percentage points lower than the multilat-

²⁰This fraction happens to be the same as the weighted mean share, i.e., the share of spending on items in *all* product categories that were also bought in January 2012.

eral index. As these figures imply, the CCDI index is much closer in value to the direct comparison fixed-base index than the chained Törnqvist index.

Panel (b) of Figure 4.3 replicates panel (a), but for a quarterly index. The quarterly CCDI and fixed based Törnqvist index resemble their monthly counterparts (though smoothing over month-to-month changes in the index). The extent to which the chained Törnqvist index diverges from the CCDI and fixed base indexes is reduced at the quarterly frequency, however, it remains the case that the chained index records much lower cumulative inflation.

Figure 4.3: *Chain drift bias: CCDI vs bilateral Törnqvists*



Note: Figures show index number values for the CCDI multilateral index, the Törnqvist fixed base index and a monthly chained Törnqvist index. The indexes are calculated across all fast-moving consumer goods.

The monthly chained Törnqvist index for all items falls by 16 ppt over the whole period, an implausibly large price change. Figure A3 in the Online Appendix shows the distribution of cumulative price changes across all 178 product categories. Individual product categories exhibit even more extreme price changes using this measure. 10 product categories record price changes of less than -50% and above 20% for 9 products categories.

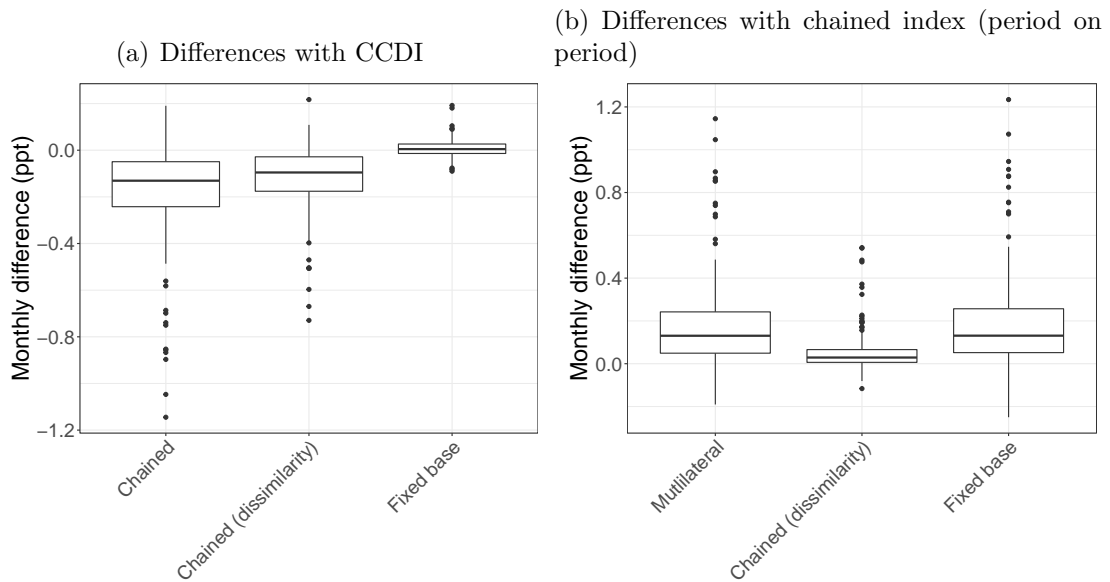
The CCDI index has an advantage over the fixed-based Törnqvist index in that it does not require all the products included in the index to be available in the base (first) period, and hence it avoids becoming as increasingly unrepresentative over time. Nonetheless, Figure 4.3 shows that over 2012–2019 the CCDI and fixed-base Törnqvist do provide a similar picture of fast-moving consumer good inflation. However, in general, this need not be the case.

Panel (a) in Figure 4.4 shows the difference in the distribution of final period values between the CCDI index and the chained and fixed-base Törnqvist indexes. We also show values for chained Törnqvist indexes, where the chaining uses the

predicted share dissimilarity method. For this and subsequent plots, the ‘boxes’ show the lower quartile, median and upper quartiles of the distribution; the lower ‘whisker’ line shows lower quartile less 1.5 times the interquartile range; and the upper whisker shows the upper quartile plus 1.5 times the interquartile range. The remaining points are outlier values.

Comparing the CCDI with the chained Törnqvist index suggests that, for most products, chain drift is negative. 22 of 178 product categories have positive chain drift bias when the Törnqvist is chained with the standard period-on-period approach.²¹ Differences between the CCDI and Törnqvist index chained using the dissimilarity index are smaller but still substantial. The median average monthly difference between the dissimilarity chained Törnqvist index is -0.09 ppt, compared to -0.13 ppt for Törnqvist index that is chained period-on-period.

Figure 4.4: *Average monthly difference between CCDI and bilateral Törnqvist indexes by product*



Note: In panel (a), each box plot summarizes the distribution (across product categories) of differences in average monthly inflation rates between the CCDI index calculated over the whole period and i) a bilateral Törnqvist chained period-on-period ii) a bilateral Törnqvist chained with using the predicted share dissimilarity approach and iii) a fixed-base Törnqvist. In panel (b), each box plot summarizes the distribution of differences in average monthly inflation rates between the period-on-period chained index bilateral Törnqvist, and i) a CCDI index calculated over the whole period ii) a bilateral Törnqvist chained with using the predicted share dissimilarity approach and iii) a fixed-base Törnqvist. We exclude outliers (the three products with the largest positive and three largest negative amounts of chain drift bias) from each plot.

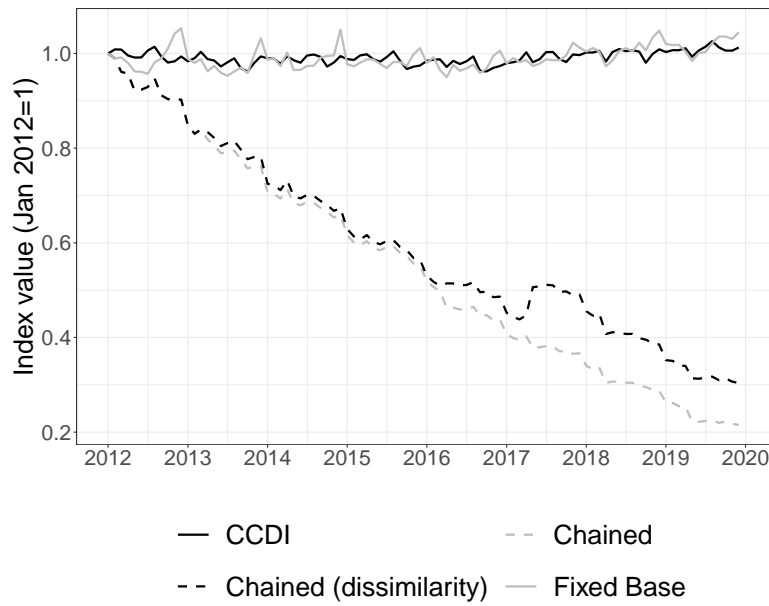
The differences between the CCDI and the fixed-base Törnqvist index are smaller than for the chained indexes, but in some cases, they are still significant. In most

²¹Prominent among these are fresh fruits. These account for three of the four product categories with the most positive chain drift bias, which are citrus fruits, apples, chilled breads, pears.

cases, these are product categories that have a high degree of product churn, reflecting bias arising from the fixed-base index becoming unrepresentative over time.

Differences between the CCDI index and the index chained using the dissimilarity approach could in principle reflect biases with either index. Panel (b) of Figure 4.4 shows the differences between the period-on-period chained Törnqvist index and the other approaches. This shows that the dissimilarity approach to chaining tends to yield results that are mostly very similar to chaining period-on-period, which as we have seen tends to imply unrealistic price changes.

Figure 4.5: *Monthly CCDI and bilateral Törnqvist indexes for Chocolate and confectionery*



Note: Figure shows index number values for the CCDI multilateral index calculated over the whole period, the Törnqvist fixed base index, a monthly chained Törnqvist index chained period-on-period, and a Törnqvist index chained using the dissimilarity approach for the product Chocolate and Confectionery.

Closer inspection of how this approach works in practice suggests that, in many cases, the most similar period according to the dissimilarity measure is the immediately preceding period. Consequently, it is often similar to the chained period-on-period index. To illustrate this aspect of the dissimilarity index, we show the time paths of the different indexes for a particular product category – Chocolate and Confectionery – in Figure 4.5. This category suffers from particularly high chain drift bias (with the second highest difference between the CCDI index, or the fixed base index, and the chained Törnqvist) and is also one for which the dissimilarity index offers one of the greatest reductions in chain drift bias relative to the period-on-period approach. Despite this, the reduction in chain drift bias when using the dissimilarity chaining approach is limited. The index chained using the dissimi-

larity method largely tracks the value of the Törnqvist chained period-on-period from 2012-2016, as the two indexes both chain on the same period. This continues until 2016, when a different chaining period is selected and the two indexes start to diverge. Figure 4.5 also highlights that, while the fixed-base index yields a similar final value to the CCDI, its path over time is considerably more volatile. This reflects the high seasonal churn in this product category.

4.1 Different multilateral indexes

In this section we quantify the difference in measured inflation across the four multilateral price indexes discussed in Section 2.3 – the CCDI, GEKS-Fisher, GEKS-Walsh and GK indexes. In each case we compute the price index, for each product category, calculated over the whole period, i.e. using all 96 year-months of the data. As this does not entail any linking, the resulting indexes do not suffer from chain drift. To summarize the difference in inflation rates, for each of the GEKS-Fisher, GEKS-Walsh, and GK indexes, we compute the difference in average monthly price changes up to the final period of data (December 2019) with that obtained with the CCDI index.²² This yields three distributions of inflation differences (across product categories), which we summarize as box plots in Figure 4.6.

The differences between the GEKS-Fisher and GEKS-Walsh indexes and the CCDI index in the final period of data are small for almost all product categories, which is perhaps unsurprising as the bilateral versions of these indexes approximate each other to the second order. However, for a small number of product categories, the difference is large. The GEKS-Fisher index records an average monthly inflation rate 0.5 ppt smaller for chilled flavoured milk than the CCDI index, and an average inflation rate 0.11 ppt smaller for ambient cakes and pastries. Cumulating these differences across all 96 months imply differences in the final period indexes of 31 ppt and 13 ppt, respectively. The GEKS-Walsh index also gives a lower price change than the CCDI for these items, yielding a 0.22 ppt smaller average monthly inflation rate for chilled flavoured milk and 0.02 ppt smaller inflation rate for ambient cakes and pastries. At the same time, the GEKS-Fisher index records an average monthly inflation rate that is 0.11 ppt greater than the CCDI index for other vegetables, while the GEKS-Walsh records an average monthly price change that is 0.02 ppt greater than the CCDI index for this product category.


These occasional differences between the CCDI and other GEKS indexes appear to reflect anomalously large price and quantity changes for particular items that

²²The average monthly inflation rate for a given index over the 96 months we consider is calculated as the difference in $x^{\frac{1}{96}} - 1$ where x is the final period value of the index.

occur in a single month, which then have a persistent impact on the cumulative index. Because the expenditure shares for these products change less than their associated quantities and prices in these cases, the CCDI is less affected than the (quantity-based) GEKS Fisher and Walsh indexes. This suggests that the CCDI can be less sensitive to certain measurement errors.²³

The differences between the average monthly inflation rates calculated with the GK and CCDI indexes are typically larger than they are for either the GEKS-Fisher or GEKS-Walsh indexes; the 25th and 75th percentiles of differences are -0.01 ppt and 0.02 ppt (compared with -0.002 and 0.002 ppt between GEKS-Fisher and CCDI index, and -0.004 and 0.003 ppt between the GEKS-Walsh and CCDI index). The GK index also records a much *greater* price increase for chilled flavoured milk than the CCDI index (unlike the other indexes, which are both substantially smaller than the CCDI). The GK index records an average monthly inflation rate that is 0.2 ppt greater for this item than the CCDI index.

Figure 4.6: *Average monthly inflation rates relative to CCDI index*



Note: Each box plot summarizes the distribution (across product categories) of differences in average monthly inflation rates between the index named in the horizontal axis and the CCDI index. Indexes are calculated using all 96 year-months of data.

²³For example: in the case of chilled flavored milk, in a single month the CCDI index falls by 2.5% while the GEKS Walsh index falls by 7% and the GEKS Fisher by an implausible 43%. This accounts for most of their final period differences. Closer inspection reveals this to be a case where there was a sharp change in observed quantities for three products (perhaps due to a change in the units of measurement being used). Expenditure levels remained relatively constant, implying large price changes. As a share weighted index, the CCDI index was much less sensitive to these changes than the quantity-weighted GEKS Fisher and GEKS Walsh indexes.

Our results point towards the CCDI index being preferred to the other multilateral indexes for measuring inflation with high frequency data. The advantage of the CCDI over the GK index is most substantial – while CCDI and the GEKS indexes mostly agree, the GK index yields divergent results. This is perhaps unsurprising: unlike the other indexes it is not based on an underlying superlative bilateral index but rather is consistent with unrealistic assumptions on underlying consumer preferences. The advantage of CCDI over the GEKS indexes is less clear-cut, but our results suggest that the former is less sensitive to large outliers likely to be driven by measurement error. In practice, NSIs are unable to use multilateral indexes calculated over all periods for headline CPIs, as this would require historical revisions when new data periods arrive. Hence, next we consider rolling-window versions of these indexes, which depend not only on the choice of index number formula but also the linking method, i.e. the method used to splice together the rolling windows.

5 Chain drift bias in spliced indexes

In this section, we quantify the chain drift bias that results from implementing the different linking procedures for extending multilateral indexes that we discuss in Section 2.4. We also show how chain drift bias changes with different window lengths used in the rolling windows that are spliced together. We do this by comparing average monthly inflation rates from spliced indexes, for a given multilateral index number, with the non-spliced series computed using all 96 year-months of data. As the latter satisfies the multiperiod identity test and is hence free of chain drift, this provides a direct measure of chain drift bias in the spliced series. We undertake this comparison for each product category for each of the CCDI, GEKS-Fisher, GEKS-Walsh, and GK index numbers.

5.1 Splicing methods

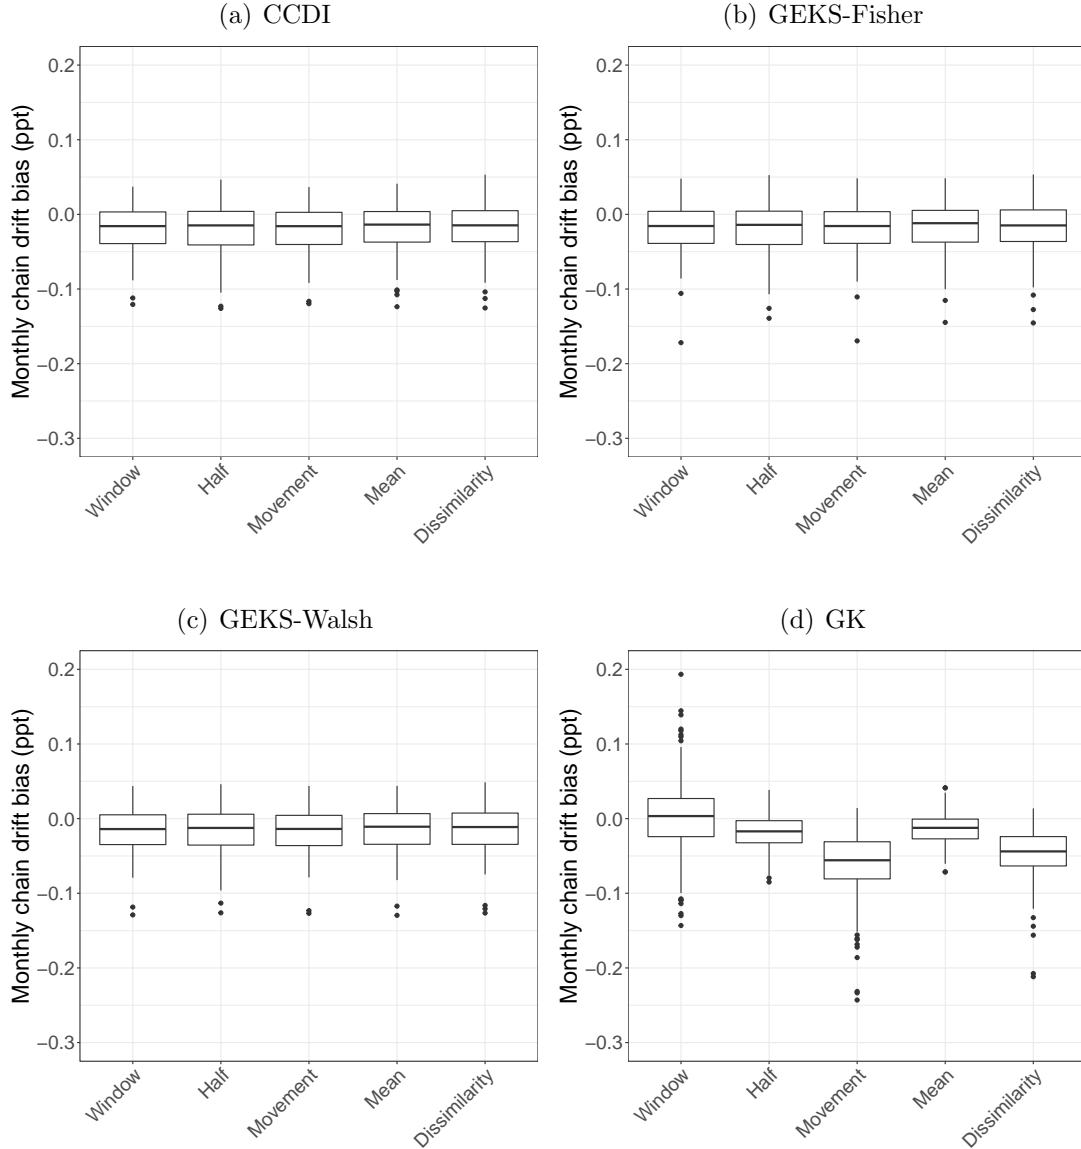
We first hold the window length fixed at 25 months and compare different linking methods: the window, half, movement and mean splices, and using the predicted share measure of relative price dissimilarity to select the splicing period.²⁴

In Figure 5.1 we show boxplots that summarize the distribution of differences between spliced indexes, and the non-spliced index using all periods of data, across product categories. Each of the four panels corresponds to a different multilateral

²⁴Additional results using the half-splice on the published series, fixed-base expanding window and fixed-base moving window are reported alongside the mean splice in the Online Appendix.

index number, and each boxplot within a panel corresponds to a different splicing method. A few individual product categories exhibit a very high degree of chain drift bias. To make the plots easier to read, we remove these extreme cases. Specifically, we remove the three product categories with the largest positive values of chain drift, and the three product categories with the largest negative values, in each plot.

Figure 5.1: *Chain drift bias with different splicing methods (25 month window)*



Note: Each box plot summarizes the distribution (across product categories) of differences in average monthly inflation between the spliced index (over a 25 month window) using the linking method named in the horizontal axis and the corresponding non-spliced index. We exclude the products with the three largest positive and negative values for chain drift in each plot.

Panel (a) shows that the distribution of chain drift bias is relatively stable across different splicing methods for the CCDI index. For each method, just under three quarters of product categories exhibit negative chain drift bias, with the remainder exhibiting positive bias. The median bias ranges from -0.02 ppt with the movement splice to -0.01 ppt with the mean and half splices, and the interquartile range of biases ranges from 0.04 to 0.05 ppt. The median chain drift bias with the mean splice implies a cumulative 1.3 ppt difference between the spliced and non-spliced index by the final period.

Panels (b) and (c) show that results for the GEKS-Walsh and GEKS-Fisher indexes are similar to those for the CCDI index. However, extreme cases of chain drift bias appear more common when we calculate prices using the GEKS-Fisher index.

Panel (d) shows that the GK index is more sensitive to the linking method than the CCDI or GEKS indexes. In particular, the movement splice results in a larger degree of chain drift bias for the GK index than for the other multilateral index numbers. The median value of the average monthly chain drift bias for the movement splice is -0.06 ppt, which is much greater than for other splicing methods. The window splice also leads to larger dispersion in rates of chain drift bias under the GK index than for other indexes. In addition, the most extreme outliers (not shown in the plots) are much higher for the GK index than for the other indexes. For instance, for the GK, seasonal biscuits have a chain drift bias of -0.99 ppt per month with the window splice, and -0.64 ppt with the movement splice.²⁵

Figure 5.1 provides the first empirical evidence on the use of the predicted share dissimilarity method for splicing rolling window multilateral indexes. At least in the first three panels, the performance is very similar to the other methods. These results are more encouraging than for its use in chaining bilateral indexes; see Figure 4.5 in Section 4.

Overall our findings suggest the CCDI and GEKS indexes are robust to the choice of extension method, in the sense that the distribution of chain drift bias is similar across them. Chain drift bias for GK index is more sensitive to the extension method, with the half and mean splice producing the best results. The mean splice avoids the risk of linking solely on a period which may happen to exhibit unusual spending and price patterns, which is potentially problematic for the window, half

²⁵The index for seasonal biscuits also shows a lot of chain drift when calculated using the CCDI index for example, but in this case it is much smaller: -0.05 ppt when calculated using the window splice and -0.06 ppt when calculated with the movement splice. Using the mean or half splice mitigates somewhat the extreme chain drift bias for this item with the GK index. The chain drift bias is -0.28 ppt with the mean splice and -0.16 ppt with the half splice.

and movement splicing methods. For predicted share dissimilarity method, the linking period is not known *ex ante* and will change between pairs of windows being linked, making it hard to explain to users. Taken together with the empirical evidence in Figure 5.1, the mean splice then seems to be the preferred choice.

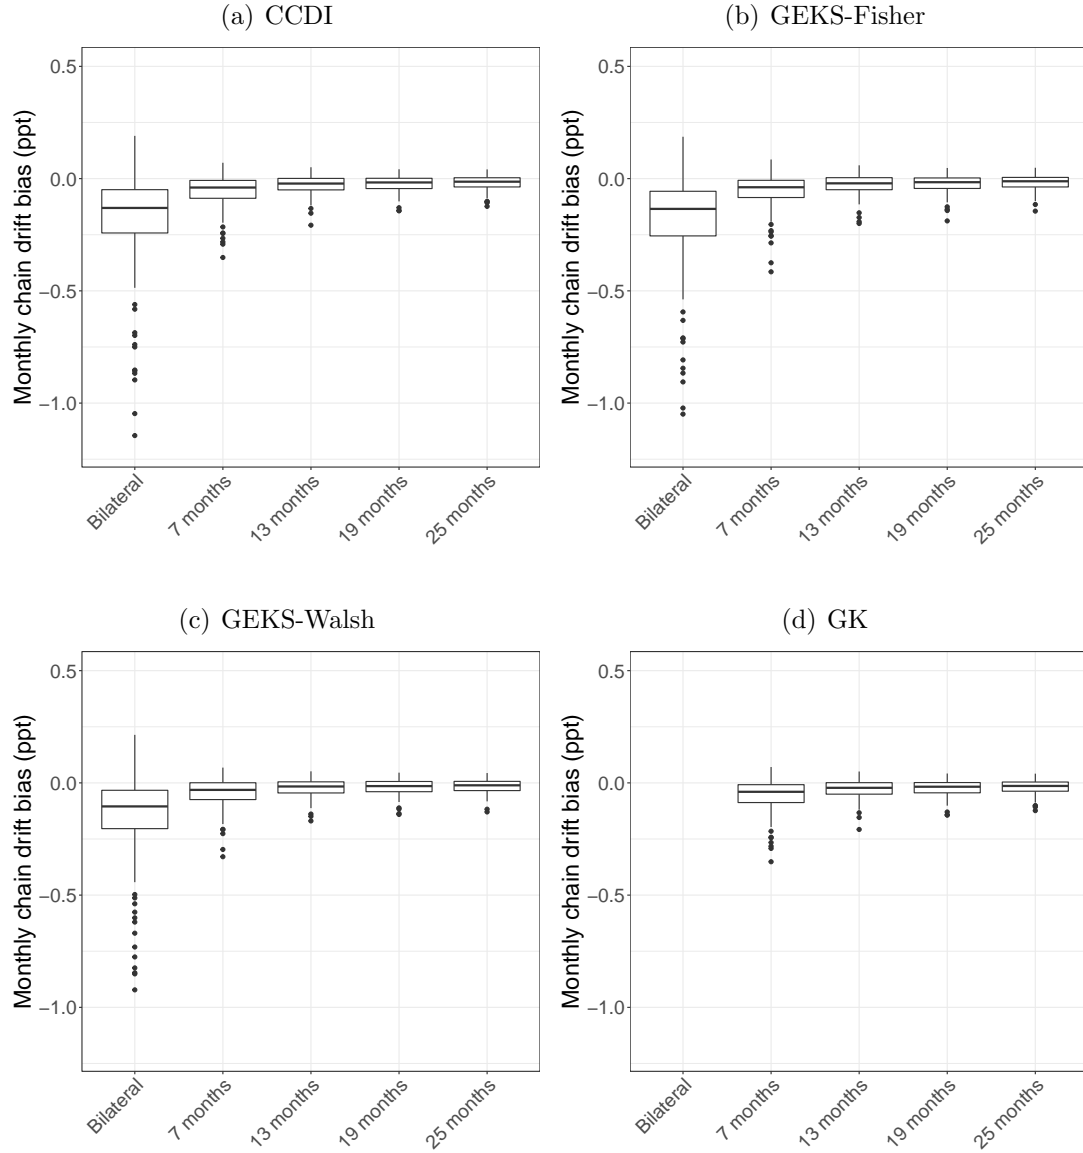
5.2 Different window lengths

In Figure 5.2 we summarize the impact of different window lengths on average monthly chain drift bias, in all cases using the mean splice. The figure is structured similarly to Figure 5.1 – each panel represents a different index number, and within each panel, the boxplots correspond to different window lengths. As before, we exclude the products with the three largest positive and negative values for chain drift in each case. In the case of the CCDI, GEKS-Fisher and GEKS-Walsh, we also include the difference between the multilateral index computed over the full period and their corresponding bilateral indexes; these are equivalent to calculating the multilateral index with a window length of one month.

For all index numbers, longer windows lengths lead to considerably less chain drift bias. For a 25 month window (the longest we consider), the distribution of chain drift bias, under CCDI, GEKS-Fisher, GEKS-Walsh and GK are similar. With a 25 month window, median average monthly chain drift bias is -0.01 ppt for all indexes and the interquartile range is 0.04 ppt for the CCDI index, GEKS-Walsh, GEKS-Fisher and 0.03 ppt for the GK index. In contrast, for a 13 month window, under the CCDI index, the median chain drift bias is -0.02 ppt and the interquartile range is 0.05 ppt. Cumulating across all months, the median bias from using a 13 month window with the CCDI index would be 1.7ppt compared to 1.3 ppt when using a 25 month window.

Figure 5.2 also demonstrates that, while spliced indexes can exhibit chain drift bias even with the longest window lengths we consider, it is noticeable that even short window lengths perform considerably better than bilateral indexes. The median average monthly chain drift bias for the bilateral Törnqvist implied by this measure is -0.13 ppt, around 10 times greater than the bias with a 25 month window length.

Figure 5.2: *Chain drift bias using different window lengths (using mean splice)*



Note: Each box plot summarizes the distribution (across product categories) of differences in average monthly inflation between the spliced index (using the mean spliced) computed over the window length named in the horizontal axis and the corresponding non-spliced index. We exclude the products with the three largest positive and negative values for chain drift in each plot. In the case of the CCDI, GEKS-Fisher and GEKS-Walsh, we also include the chain drift bias associated with their corresponding bilateral indexes (equivalent to using a window length of one month). The GK index does not have a corresponding bilateral index.

5.3 Drivers of chain drift bias

The results from the previous section highlight the importance of using relatively long window lengths for spliced multilateral index numbers to minimize the degree

of chain drift bias. Chain drift biases decline noticeably from 7 months to 25 months window lengths across all index number methods.

Here we assess under what circumstances chain drift bias is likely be a significant issue. In particular, we consider how chain drift bias at different window lengths relates to five possible drivers (all measured separately for each product category)

- **Monthly churn:** This is measured by the share of spending on products in the current month that were not observed being purchased in the previous month.
- **Annual churn:** This is measured by the share of spending on products in the current year that were not observed being purchased in the previous year. Run-out sales at the end of product life-cycles have been identified as a potentially important cause of chain drift (Melser and Webster (2021)).
- **Seasonality in pricing (‘weak seasonality’):** We measure this by, for each product category, estimating a regression of log price on product fixed effects and month dummies. We measure the degree of seasonality in a product category as the difference between the largest and smallest month dummy from this regression.
- **The frequency of price promotions:** This is the percentage of transactions each year that are observed with price promotions.
- **The frequency of quantity promotions:** This is the percentage of transactions each year that are observed with quantity promotions (for example, two-for-one offers).

Table 5.1 shows the distribution of these measures across product categories. Monthly churn is highest for the category seasonal biscuits; on average, 14.6% of spending each month were on products not observed purchased in the previous month. It also exhibits the greatest annual rates of churn and seasonal pricing of any product category. Annual churn is also high for products like chocolate and air fresheners, while monthly churn is higher for seasonal products like vitamins, minerals and skincare. Apart from seasonal biscuits, seasonal pricing is most important for soft fruits and fortified wines. Price promotions are most important for ‘mini portions’ of dairy products and healthy biscuits, while quantity promotions are most common for fresh pasta and chilled processed poultry.

We assess the role these factors play in driving chain drift bias by undertaking the following analysis. For each multilateral index number, we compute the absolute

Table 5.1: *Summary Statistics*

Variable	Mean	Min	Pctl. 25	Pctl. 75	Max
Monthly churn	2.49	0.24	1.24	3.15	14.55
Annual churn	8.95	1.28	5.32	11.96	38.64
Seasonal pricing	7.47	2.71	5.04	8.68	27.99
Price promotions	23.54	2.22	15.85	29.89	50.27
Quantity promotions	9.63	0.029	3.90	13.91	29.75

Note: Numbers for monthly and annual churn are % of spending. Numbers for seasonal pricing are the maximum difference in average log-price between calendar quarters (conditional on product fixed effects). Numbers for price and quantity promotions are share of transactions. Summary statistics are across product categories.

value of the cumulative chain drift bias over all 96 months for a spliced series (using the mean splice) computed over both a 7 and 25 month window. We regress this variable on each of the potential drivers of chain drift bias summarized in Table 5.1. In each regression, an observation is a product category (note, we drop the three product categories with the highest degree of chain drift bias). Table 5.2 shows results using a short window length (7 months) and Table 5.3 shows results using a longer window length (25 months).

Using the 7 month window, higher rates of annual product churn are positively and statistically significantly associated with chain drift bias. This is true for all index numbers, and the effects are large. Each percentage point increase in the annual product churn measure is associated with an increase in the absolute value of cumulative chain drift bias of between 0.46 and 0.64 ppts (depending on the index). There are also strong effects associated with seasonal pricing. A one log point increase in the difference between peak and trough prices within each year is associated with an increase in cumulative chain drift bias of between 0.39 and 0.48 ppt.

Using a 25 month window length however, the effects of seasonal pricing are small and no longer statistically significant. Rather, it is higher rates of product churn that are the main determinant of chain drift bias. Each percentage point increase in annual churn is associated with between a 0.13 and 0.15 ppt increase in chain drift for the CCDI, GEKS-Fisher and GEKS-Walsh indexes. This is large relative to the overall bias at a window length of 25 months. The notable exception to this pattern is the GK index, which appears more sensitive to differences in monthly than annual churn. Each percentage point increase in annual churn leads to a 0.06 ppt increase in chain drift bias for the GK, while each percentage increase in monthly churn leads to an increase in chain drift bias of 0.14 ppt.

The results suggest that product churn is a key determinant of chain drift bias. Monthly and annual churn contribute to 88% of the total explained variance in chain drift in the 25 month spliced CCDI index. Longer window lengths help to mitigate the effects of annual churn for the CCDI and GEKS indexes, and the effects of seasonality in pricing. High-frequency (monthly) churn appears to be a particular problem for the GK index.

Table 5.2: *Determinants of chain drift bias (7 month window length)*

	CCDI	GEKS-Fisher	GEKS-Walsh	GK
	(1)	(2)	(3)	(4)
Monthly churn	-0.023 (0.110)	0.005 (0.117)	0.016 (0.095)	0.015 (0.099)
Annual churn	0.495** (0.243)	0.477* (0.257)	0.457** (0.211)	0.639*** (0.217)
Pricing seasonality	0.479*** (0.107)	0.478*** (0.113)	0.424*** (0.092)	0.389*** (0.095)
Price promotions	-0.105* (0.059)	-0.099 (0.063)	-0.086* (0.051)	-0.125** (0.053)
Quantity promotions	0.030 (0.041)	0.031 (0.043)	0.019 (0.035)	0.024 (0.037)
Observations	175	175	175	175
R ²	0.202	0.181	0.213	0.220

Note: All indexes are extended using the mean splice. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Table 5.3: *Determinants of chain drift bias (25 month window length)*

	CCDI	GEKS-Fisher	GEKS-Walsh	GK
	(1)	(2)	(3)	(4)
Monthly churn	0.074 (0.108)	0.076 (0.122)	0.021 (0.103)	0.139* (0.074)
Annual churn	0.149*** (0.046)	0.154*** (0.052)	0.129*** (0.044)	0.061* (0.033)
Pricing seasonality	0.001 (0.048)	0.023 (0.053)	0.001 (0.045)	-0.028 (0.033)
Price promotions	-0.009 (0.027)	-0.021 (0.030)	-0.006 (0.026)	-0.023 (0.018)
Quantity promotions	-0.024 (0.019)	-0.038* (0.021)	-0.031* (0.018)	-0.001 (0.012)
Observations	175	175	175	175
R ²	0.110	0.100	0.085	0.073

Note: All indexes are extended using the mean splice. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

6 Conclusions

Accurate and timely inflation measurement is important for a wide range of economic policies, particularly in times of high and volatile inflation. The increasing availability of transaction level datasets means it is now possible to measure inflation using high frequency data sources, meaning potentially faster data collection, compiling and release of the CPI using far more information than traditionally has been the case. Various methods for measuring inflation with such data have been proposed and even implemented by NSIs, but to date there remains a paucity of evidence systematically comparing them. Such empirical comparisons are important, as each method can result in chain drift bias, which can be substantial.

In this paper we help fill this evidence gap by providing a systematic quantitative comparisons between the competing methods for measuring month-to-month inflation using long-run transaction (scanner) data that covers the 178 product categories that comprise the fast-moving consumer good segment of the economy. These comparisons include proposed methods that have had until now limited empirical applications or none at all.

We find that the CCDI multilateral index, with a 25 month rolling window and a mean splice for linking, is the preferred approach. We also provide evidence on the determinants of chain drift bias for spliced multilateral indexes, which we believe is novel to the literature. We find that product churn (i.e., entry and exit) is a key determinant of chain drift bias.

There are several potential avenues for future research in this area. One outstanding question concerns variation in chain drift bias across months within a price series. In our analysis, we focus on an average measure of monthly chain drift bias over a long time period, but the variance of the bias across months (for a given average bias) and how this varies across low and high inflation periods also impacts the reliability of the index as a real-time inflation measure. In addition, future work could assess the impact of disappearing products on the behavior of multilateral indexes, as well as the use of imputation methods for temporarily missing products.

References

- Australian Bureau of Statistics (2017). An implementation plan to maximise the use of transactions data in the CPI. Information Paper 6401.0.60.004, Canberra, Australia.
- Balk, B. M. (1981). A simple method for constructing price indices for seasonal commodities. *Statistische Hefte* 22(1), 72–78.

- Broda, C. and D. E. Weinstein (2010). Product Creation and Destruction: Evidence and Price Implications. *American Economic Review* 100(3), 691–723.
- Caves, D. W., L. R. Christensen, and W. E. Diewert (1982). Multilateral Comparisons of Output, Input, and Productivity Using Superlative Index Numbers. *The Economic Journal* 92(365), 73–86. Publisher: [Royal Economic Society, Wiley].
- Chessa, A. G. (2021). Extension of multilateral index series over time: Analysis and comparison of methods. *Department of Consumer Price Statistics Netherlands*.
- Coibion, O., Y. Gorodnichenko, and G. H. Hong (2015). The Cyclicalities of Sales, Regular and Effective Prices: Business Cycle and Policy Implications. *American Economic Review* 105(3), 993–1029.
- Coibion, O., Y. Gorodnichenko, and R. Kamdar (2018). The Formation of Expectations, Inflation, and the Phillips Curve. *Journal of Economic Literature* 56(4), 1447–1491.
- de Haan, J. (2015). A Framework for Large Scale Use of Scanner Data in the Dutch CPI. *Presented at the 14th Meeting of the Ottawa Group, Tokyo*.
- de Haan, J. and H. A. van der Grient (2011). Eliminating Chain Drift in Price Indexes Based on Scanner Data. *Journal of Econometrics*.
- Diewert, W. E. (1976). Exact and superlative index numbers. *Journal of Econometrics* 4(2), 115–145.
- Diewert, W. E. (2002). Similarity and dissimilarity indexes: An axiomatic approach. *University of British Columbia Discussion Paper* (02-10).
- Diewert, W. E. (2022). Scanner Data, Elementary Price Indexes and the Chain Drift Problem. In D. Chotikapanich, A. N. Rambaldi, and N. Rohde (Eds.), *Advances in Economic Measurement*, pp. 445–606. Singapore: Springer Nature Singapore.
- Diewert, W. E., Y. Finkel, D. Sayag, and G. White (2022). Chapter 9: Seasonal products. In *Consumer Price Index Theory*. International Monetary Fund.
- Diewert, W. E. and K. J. Fox (2022). Substitution Bias in Multilateral Methods for CPI Construction. *Journal of Business & Economic Statistics* 40, 355–369.
- Diewert, W. E., K. J. Fox, and J. de Haan (2016). A newly identified source of potential cpi bias: Weekly versus monthly unit value price indexes. *Economics Letters* 141, 169–172.
- Eichenbaum, M., N. Jaimovich, and S. Rebelo (2011). Reference Prices, Costs, and Nominal Rigidities. *American Economic Review* 101(1), 234–262.
- Eltető, O. and P. Köves (1964). On a Problem of Index Number Computation Relating to International Comparisons. *Statistikai* 42(10), 507–518.
- Geary, R. C. (1958). A Note on the Comparison of Exchange Rates and Purchasing Power Between Countries. *Journal of the Royal Statistical Society. Series A (General)* 121(1), 97–99. Publisher: [Royal Statistical Society, Wiley].
- Gini, C. (1931). On the circular test of index numbers. *Metron* 9(9), 3–24.
- Handbury, J., T. Watanabe, and D. E. Weinstein (2013). How Much Do Official Price Indexes Tell Us About Inflation? Working Paper 19504, NBER, Cambridge, MA.
- Hill, R. J. (1999, 02). Comparing Price Levels across Countries Using Minimum-Spanning Trees. *The Review of Economics and Statistics* 81(1), 135–142.

- Hill, R. J. (2001). Measuring inflation and growth using spanning trees. *International Economic Review* 42(1), 167–185.
- ILO, IMF, OECD, UNECE, Eurostat, and The World Bank (2004). *Consumer price index manual: Theory and practice*. Geneva: International Labour Office: P. Hill.
- Inklaar, R. and W. E. Diewert (2016). Measuring industry productivity and cross-country convergence. *Journal of Econometrics* 191(2), 426–433.
- Ivancic, L., W. Erwin Diewert, and K. J. Fox (2011). Scanner data, time aggregation and the construction of price indexes. *Journal of Econometrics* 161(1), 24–35.
- Jaravel, X. (2019). The Unequal Gains from Product Innovations: Evidence from the U.S. Retail Sector. *The Quarterly Journal of Economics* 134(2), 715–783.
- Jaravel, X. and M. O’Connell (2020). Real-time price indices: Inflation spike and falling product variety during the Great Lockdown. *Journal of Public Economics* 191, 104–270.
- Kaplan, G. and S. Schulhofer-Wohl (2017). Inflation at the Household Level. *Journal of Monetary Economics* 91, 19–38.
- Khamis, S. H. (1970). Properties and Conditions for the Existence of a New Type of Index Number. *Sankhya B* 32, 81–98.
- Khamis, S. H. (1972). A New System of Index Numbers for National and International Purposes. *Journal of the Royal Statistical Society: Series A (General)* 135(1), 96–121. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.2307/2345041>.
- Könus, A. A. (1924). The Problem of the True Index of the Cost or Living (translated in 1939). *Econometrica* 7, 10–29.
- Krsinich, E. (2016). The FEWS Index: Fixed Effects With a Window Splice. *Journal of Official Statistics* 32, 375–404.
- Lowe, J. (1823). *The Present State of England in Regard to Agriculture, Trade and Finance, second ed.* Longman, Hurst, Rees, Orme and Brown.
- Melser, D. and M. Webster (2021). Multilateral methods, substitution bias, and chain drift: Some empirical comparisons. *Review of Income and Wealth* 67(3), 759–785.
- Nakamura, E. and J. Steinsson (2013). Price Rigidity: Microeconomic Evidence and Macroeconomic Implications. *Annual Review of Economics* 5(1), 133–163.
- Reinsdorf, M. B. (1998). Formula Bias And Within-Stratum Substitution Bias in the U.S. CPI. *Review of Economic and Statistics* LXXX, 175–187.
- Szulc, B. (1964). Indices for Multiregional Comparisons. *Przegląd Statystyczny* 3, 239–254.

APPENDIX: FOR ONLINE PUBLICATION

Inflation measurement with high frequency data

Kevin J. Fox, Peter Levell, Martin O’Connell
October, 2023

A Alternative splicing approaches

In this appendix we describe and evaluate several alternative splicing methods (in particular splicing on the published series, the fixed-base moving window and fixed-base expanding window) that have been proposed in the literature.

Slicing on the published series. An alternative to the rolling window splice is to splice on the published series. In particular, when a new data point in period $t = s + T$ becomes available, the corresponding new sequence $\mathbb{P} = (\mathbb{P}^{t-T+1}, \dots, \mathbb{P}^t)$ can be directly spliced onto the published series $\rho_1, \dots, \rho_{t-1}$. Let τ be the link period (e.g., the movement, window, or half splice), then the price level for period $t + 1$ is given by:

$$\rho_{t+1}(\tau) = \rho_{\tau+s}(\tau) \frac{\mathbb{P}^{t+1}}{\mathbb{P}_{\tau+s}}$$

Chessa (2021) suggests implementing this method with the half splice.

The fixed-base moving window. An alternative to the rolling window approach is to link the current price level, computed on the most recent multilateral sequence, with the spliced series based on a period that is fixed in calendar (rather than relative) time. This is known as the fixed-base moving window (FBMW). Denote by \mathfrak{t} a link period defined in calendar time. Under the FBMW, the spliced price level in period t is given by:

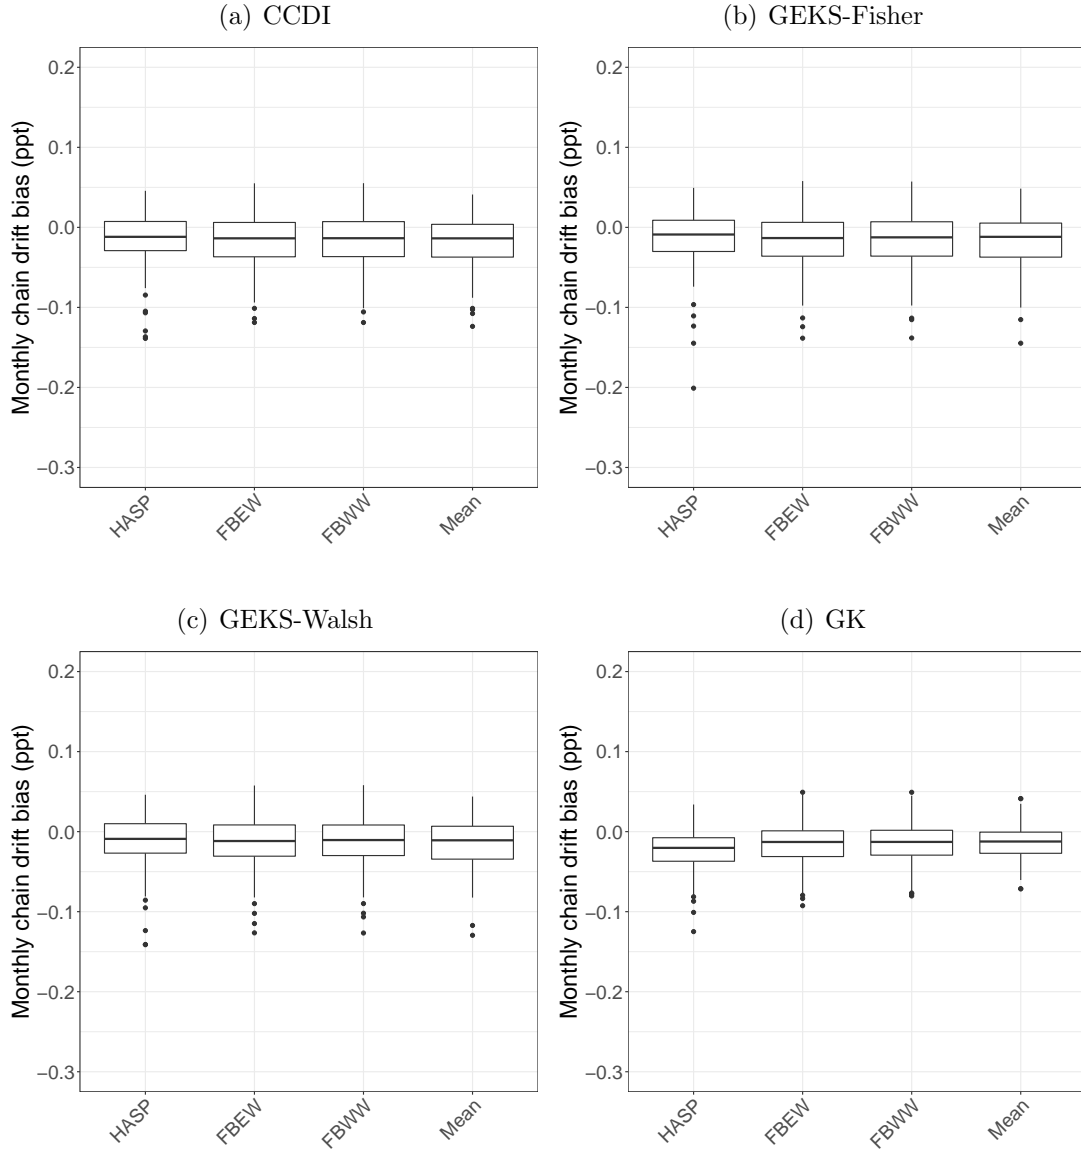
$$\rho_t = \rho_{\mathfrak{t}} \frac{\mathbb{P}_t^t}{\mathbb{P}_t^{\mathfrak{t}}},$$

The fixed-base expanding window. The fixed-base expanding window (FBEW), is similar to the fixed-based moving window, but rather than each multilateral sequence being of T periods long, it *expands* the window each period to include the latest period of data.

For example, with monthly data, and a December base month, the window used to compute the new data point in January includes only December and January. In February, it will include December, January and February, and so on until it includes all months in a given window.

Figure A1 plots the distribution of monthly chain drift biases associated with the half-splice on the published series, FBEW and FBMW approaches.

Figure A1: *Chain drift bias with different splicing methods (25 month window)*



Note: Each box plot summarizes the distribution (across product categories) of differences in average monthly inflation between the spliced index (over a 25 month window) using the linking method named in the horizontal axis and the corresponding non-spliced index. We exclude the products with the three largest positive and negative values for chain drift in each plot.

B Additional Figures

Table A1: *Product categories (1)*

	Annual spending share (%)		
	mean	min	max
Bakery			
Ambient Cakes and Pastries	1.45	1.40	1.52
Morning Goods	1.62	1.49	1.76
Total Bread	1.74	1.59	2.04
Chilled Breads	0.16	0.14	0.17
Chilled Cakes	0.33	0.29	0.35
Dairy			
Butter	0.85	0.74	1.02
Cheddar	1.59	1.51	1.70
Continental Ex.Blue	0.49	0.39	0.56
Eggs	0.84	0.78	0.88
Fresh Cream	0.34	0.32	0.38
Fromage Fraiss	0.28	0.16	0.34
Margarine	0.64	0.52	0.85
Mini Portions	0.13	0.11	0.16
Semi-skimmed milk	1.42	1.33	1.50
Skimmed milk	0.37	0.35	0.40
Territorials Ex.Blue	0.23	0.21	0.25
Total Milk	0.51	0.48	0.60
Total Processed	0.34	0.33	0.36
Total Soft White	0.24	0.23	0.25
Whole milk	0.64	0.59	0.71
Yoghurt	1.63	1.58	1.65
Yoghurt Drinks And Juices	0.29	0.28	0.32
Fresh fruit and vegetables			
Apples	0.84	0.80	0.86
Bananas	0.60	0.57	0.63
Brassicas	0.61	0.58	0.67
Chilled Prepared Fruit and Veg	0.95	0.82	1.05
Citrus	0.76	0.71	0.83
Legumes	0.21	0.18	0.23
Nuts - fruit	0.22	0.14	0.28
Other Vegetables	0.89	0.80	0.95
Pears	0.22	0.19	0.24
Potatoes	1.20	1.02	1.48
Root Crops	0.84	0.76	0.95
Salads	1.78	1.66	1.91
Soft Fruit	2.25	1.88	2.55
Tropical Fruits	0.49	0.40	0.55

Table A2: *Product categories (2)*

	Annual spending share (%)		
	mean	min	max
Fresh meat and fish			
Chilled Prepared Fish	0.20	0.17	0.24
Shellfish	0.19	0.18	0.22
Wet/Smoked Fish	0.87	0.75	0.97
Chilled Burgers and Grills	0.28	0.23	0.32
Chld Frnkfurter/Cont Ssgs	0.15	0.13	0.16
Fresh Bacon Joint	0.25	0.23	0.26
Fresh Bacon Rashers	0.95	0.87	1.03
Fresh Bacon Steaks	0.13	0.12	0.15
Fresh Beef	2.17	2.07	2.26
Fresh Flavoured Meats	0.15	0.13	0.16
Fresh Lamb	0.53	0.48	0.57
Fresh Pork	0.78	0.67	0.88
Fresh Sausages	0.70	0.68	0.74
Chilled Processed Poultry	0.38	0.32	0.43
Cooked Poultry	0.51	0.48	0.54
Fresh Poultry	2.31	2.26	2.34
Chilled prepared			
Chilled Desserts	0.69	0.66	0.71
Chilled Dips	0.18	0.14	0.22
Chilled Pizza and Bases	0.52	0.47	0.55
Chilled Prepared Salad	0.32	0.28	0.35
Chilled Ready Meals	2.53	2.21	2.77
Chld Sandwich Fillers	0.13	0.12	0.15
Cooked Meats	2.34	2.24	2.47
Fresh Pasta	0.16	0.13	0.17
Fresh Soup	0.11	0.10	0.12
Other Chilled Convenience	0.28	0.21	0.31
Fresh Meat and Veg Pastry	0.97	0.89	1.01
Frozen meat			
Frozen Fish	0.95	0.90	0.99
Frozen Sausages	0.11	0.09	0.11
Frozen Poultry	0.39	0.33	0.44
Frozen Meat Products	0.18	0.16	0.20
Frozen Pizzas	0.57	0.51	0.63
Frozen Potato Products	0.85	0.82	0.89
Frozen Processed Poultry	0.53	0.49	0.56
Frozen Ready Meals	0.77	0.73	0.84
Frozen Savoury Bakery	0.22	0.21	0.23
Frozen Vegetables	0.59	0.58	0.62
Frozen Vegetarian Prods	0.22	0.19	0.26
Other Frozen Foods	0.16	0.15	0.18

Table A3: *Product categories (3)*

	Annual spending share (%)		
	mean	min	max
Cupboard ingredients			
Ambient Soup	0.35	0.31	0.40
Baked Bean	0.41	0.38	0.46
Canned Fish	0.56	0.54	0.59
Canned Hot Meats	0.18	0.15	0.21
Canned Pasta Products	0.12	0.10	0.15
Canned Vegetables	0.13	0.13	0.14
Cold Canned Meats	0.14	0.12	0.16
Prepared Peas and Beans	0.15	0.15	0.16
Tinned Fruit	0.17	0.16	0.18
Tomato Products	0.27	0.27	0.28
Food Drinks	0.21	0.18	0.22
Instant Coffee	0.91	0.88	0.98
Liquid/Grnd Coffee and Beans	0.39	0.26	0.48
Tea	0.54	0.49	0.61
Breakfast Cereals	1.89	1.72	2.09
Honey	0.11	0.10	0.11
Preserves	0.16	0.15	0.18
Ambnt Salad Accompanimet	0.28	0.27	0.29
Sour and Speciality Pickles	0.13	0.13	0.13
Table Sauces	0.30	0.29	0.31
Ambient Rice and Svry Noodles	0.58	0.57	0.59
Dry Pasta	0.25	0.22	0.27
Instant Hot Snacks	0.15	0.13	0.18
Packet Soup	0.13	0.10	0.15
Ambient Cooking Sauces	0.84	0.73	0.96
Cooking Oils	0.36	0.34	0.37
Ethnic Ingredients	0.22	0.20	0.24
Flour	0.12	0.10	0.14
Herbs and Spices	0.20	0.18	0.22
Meat Extract	0.40	0.39	0.42
Home Baking	0.52	0.48	0.55
Sugar	0.30	0.24	0.42
Confectionery			
Cereal and Fruit Bars	0.38	0.36	0.41
Childrens Biscuits	0.15	0.14	0.16
Chocolate Biscuit Bars	0.44	0.40	0.47
Confectionery and Other Exclusions	0.19	0.18	0.21
Crackers and Crispbreads	0.37	0.35	0.38
Everyday Biscuits	0.34	0.31	0.37
Everyday Treats	0.40	0.37	0.42
Healthier Biscuits	0.24	0.21	0.25

Table A4: *Product categories (4)*

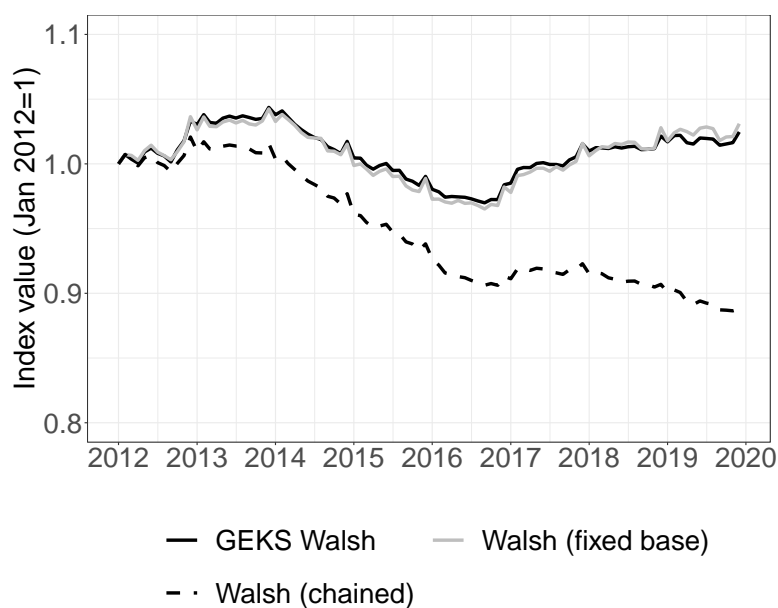
	Annual spending share (%)		
	mean	min	max
Savoury Biscuits	0.15	0.13	0.18
Seasonal Biscuits	0.13	0.12	0.14
Special Treats	0.16	0.14	0.17
Frozen Confectionery	0.35	0.32	0.40
Total Ice Cream	1.03	0.94	1.16
Chocolate Confectionery	2.47	2.38	2.60
Sugar Confectionery	0.74	0.72	0.75
Crisps	1.00	0.96	1.07
Nuts - savoury	0.26	0.23	0.28
Savoury Snacks	1.03	0.95	1.14
Drinks			
Chilled Flavoured Milk	0.12	0.11	0.13
Chilled Fruit Juice and Drink	0.71	0.63	0.79
Ambient One Shot Drinks	0.37	0.29	0.43
Ambient Fruit Yoghurt Drinks	0.39	0.30	0.51
Bottled Colas	0.56	0.52	0.61
Bottled Lemonade	0.11	0.09	0.14
Bottled Other Flavours	0.43	0.41	0.47
Canned Colas	0.53	0.49	0.60
Canned Other Flavours	0.27	0.24	0.31
Mineral Water	0.45	0.35	0.53
Total Fruit Squash	0.60	0.56	0.65
Alcohol			
Beer and Lager	1.68	1.63	1.76
Cider	0.54	0.52	0.57
Fabs	0.12	0.10	0.13
Fortified Wines	0.23	0.19	0.26
Sparkling Wine	0.46	0.33	0.54
Spirits	2.51	2.32	2.76
Wine	3.26	3.21	3.34
Household goods			
Bath and Shower Products	0.41	0.39	0.42
Deodorants	0.46	0.42	0.50
Liquid Soap	0.16	0.15	0.17
Skincare	0.58	0.55	0.64
Hair Colourants	0.18	0.15	0.20
Hair Conditioners	0.20	0.19	0.21
Shampoo	0.33	0.32	0.34
Oral Analgesics	0.20	0.19	0.23
Vitamins.Minerals/splmnts	0.35	0.32	0.37
Air Fresheners	0.36	0.33	0.39
Batteries	0.25	0.24	0.26

Table A5: *Product categories (5)*

	Annual spending share (%)		
	mean	min	max
Bin Liners	0.18	0.16	0.21
Bleaches and Lavatory Cleaners	0.28	0.27	0.30
Cleaning Accessories	0.14	0.13	0.15
Fabric Conditioners	0.43	0.36	0.47
Facial Tissues	0.25	0.25	0.26
Household Cleaners	0.42	0.41	0.43
Household Food Wraps	0.24	0.24	0.25
Kitchen Towels	0.40	0.38	0.40
Machine Wash Products	0.99	0.87	1.09
Toilet Tissues	1.29	1.25	1.34
Wash Additives	0.12	0.11	0.14
Washing Up Products	0.51	0.46	0.54
Mouthwashes	0.18	0.16	0.19
ToothPastes	0.39	0.38	0.40
Total Toothbrushes	0.20	0.18	0.21
Feminine Care	0.23	0.20	0.25
Incontinence Products	0.11	0.07	0.16
Moist Wipes	0.20	0.15	0.25
Razor Blades	0.21	0.17	0.25
Cat Litter	0.16	0.15	0.16
Cat and Dog Treats	0.55	0.47	0.62
Dog Food	0.56	0.53	0.60
Total Cat Food inc.Bulk	1.42	1.31	1.50
Total Dry Dog Food	0.14	0.11	0.16

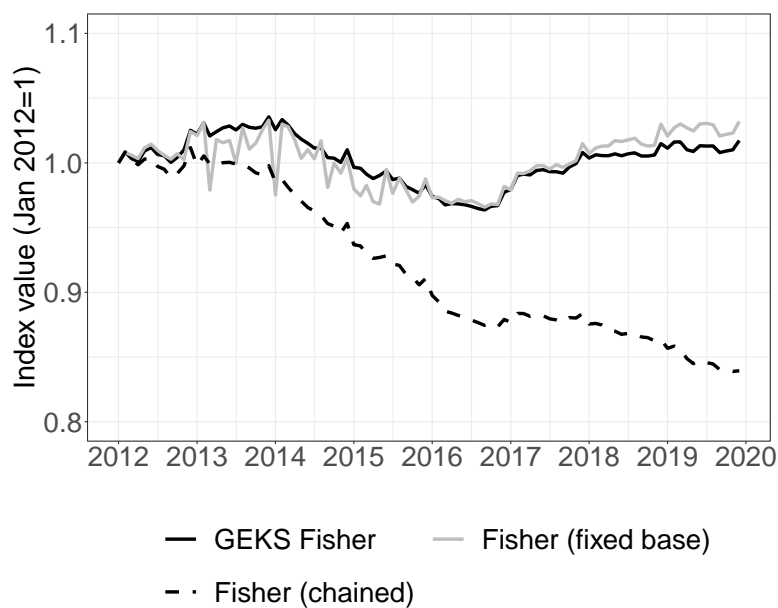
Note: Authors' calculations based on the Kantar FMCG At-Home Purchase Panel for 2012-2019.

Figure A1: *Chain drift bias: GEKS Walsh vs Bilateral Walsh*




Note: Figure shows index number values for the GEKS-Walsh multilateral index, the Walsh fixed base index and a monthly chained Walsh index. The indexes are calculated across all fast-moving consumer goods.

Figure A2: *Chain drift bias: GEKS Fisher vs Bilateral Fisher*



Note: Figure shows index number values for the GEKS-Fisher multilateral index, the Fisher fixed base index and a monthly chained Fisher index. The indexes are calculated across all fast-moving consumer goods.

Figure A3: *Distribution of cumulative inflation rates for the Törnqvist index with different chaining methods*



Note: Figure shows the distribution of cumulative price changes from January 2012 to December 2019 across 178 product categories calculated using the bilateral Törnqvist index using different monthly chaining methods.