

Inflation measurement with high-frequency data

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

July, 2025

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 and timely measures of inflation. However, there is no universally accepted method for calculating price indexes using such high-frequency data, reflecting a lack of systematic evidence on the performance of different approaches. We use a dataset covering 178 product categories, comprising all fast-moving consumer goods over eight 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

*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 accurately and in a timely way is essential for effective policymaking. Central banks, for example, rely on inflation measures when setting interest rates, and their success in maintaining price stability depends in part 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, many benefit and social insurance programs are indexed to inflation measures, meaning their capacity to protect individuals from adverse inflation shocks depends on how quickly and accurately they reflect sudden changes in consumer purchasing power.

Traditionally, National Statistical Institutes (NSIs) collect price data for consumer price indexes (CPIs) using in-person collectors. This approach yields a relatively small sample of price quotes and lacks product-level expenditure information, resulting in month-to-month changes in the index that can be noisy. In addition, index methods typically used to construct CPIs—based on historical expenditure weights at the product-aggregate level—can result in an index that becomes rapidly unrepresentative of spending patterns over time.

The increasing availability of large, comprehensive transaction-level scanner datasets, which contain near 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 key segments of the economy.¹ These datasets are now increasingly used by NSIs to produce price indexes.² However, index construction using such data presents challenges, including high rates of product entry and exit and volatile movements in prices and quantities—both of which can introduce biases into the resulting index (Ivancic et al., 2011; de Haan and van der Grient, 2011). While Diewert and Fox (2022) have recently presented theoretical and simulation

¹Scanner datasets 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).

²Retail 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 many of these agencies working on, or already incorporating, transaction data into their CPI.

evidence on the extent of substitution bases in alternative methods, empirical evidence remains limited. Moreover, there is still no universally accepted method for calculating price indexes with high-frequency transaction data.

We provide a systematic empirical comparison of alternative index number methods for computing month-to-month inflation with high-frequency transaction-level data. We use household scanner data from the Kantar FMCG At-Home Purchase Panel, in which households record purchases of fast-moving consumer goods. These include food and beverages (both alcoholic and non-alcoholic), as well as household supplies such as toiletries, cleaning products, and pet foods. The dataset has information on expenditures and transaction prices for over 290,000 unique products and more than 300 million transactions spanning eight years (2012-2019). We compute price indexes for fast-moving consumer goods as a whole and separately for each of the 178 product categories, comparing fixed-base, chained bilateral, and multilateral index methods.

We begin by showing that the use of fixed expenditure weights can cause a price index to quickly become unrepresentative: by the final month of our data, only 32% of expenditure is on products that were also purchased in the first month. Chained bilateral indexes address this by updating the basket of products over time, but they can introduce chain drift bias. This bias can arise when changes in product weights between two periods are correlated with prices in other periods (see Reinsdorf, 1998; Diewert, 2022).

We show that chained bilateral indexes can exhibit substantial chain drift. For example, a chained Törnqvist index—the type of index used by the U.S. Bureau of Labor statistics in calculating the C-CPI-U—reports cumulative inflation of -16% for all fast-moving consumer goods over 2012-2019. At the category level, it shows cumulative inflation above 20% for nine product categories and below -50% for ten product categories. Inflation of this magnitude is not plausible for this sector and time period. For indexes computed at a higher level of time aggregation, chain drift is reduced but can remain significant.

One potential solution to tackling this bias is to chain across periods that are not necessarily adjacent but are instead selected based on similarity in price structures or available products. In principle, this approach can reduce chain drift, but its empirical performance has not been extensively explored. We find that the bias associated with monthly chaining is only marginally reduced when applying this method.

A second approach to addressing 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 are free from chain drift by construction. We compare three “GEKS”-type multilateral indexes based on the Törnqvist, Fisher and Walsh indexes. Each of these underlying bilateral indexes is “superlative”, meaning it equals the true cost-of-living index defined by Könus (1924) under specific preferences that can be represented by a functional form that approximates arbitrary preferences to the accuracy of a second-order approximation (Diewert, 1976). The corresponding multilateral indexes are the Caves-Christensen-Diewert-Inklaar (CCDI), GEKS-Fisher and GEKS-Walsh indexes (where GEKS stands for Gini-Eltetö-Köves-Szulc).³

Using our dataset, the CCDI index reports cumulative inflation of 2.5% for all fast moving consumer goods. The other GEKS-type indexes generate similar results to the CCDI index. However, we find the former are more sensitive to large product-level price changes likely to reflect measurement error. NSIs might therefore prefer the CCDI index. We also evaluate the Geary-Khamis (GK) index, which is not constructed from a superlative bilateral index number. It is microfounded by a consumer model with either linear or Leontief preferences, making it more prone to substitution bias. This theoretical limitation is reflected in practice: across product categories, the 25th and 75th percentiles of the distribution of differences in average monthly inflation relative to the CCDI index are -0.002 and 0.002 percentage points (ppt) for the GEKS-Fisher index, -0.004 ppt and 0.004 ppt for the GEKS-Walsh index, and -0.01 ppt and 0.02 ppt for the GK index.

³The CCDI index is also known as the GEKS-Törnqvist index.

A practical drawback of multilateral indexes in their pure form is that they require revisions to historical numbers whenever new data becomes available, making them unsuitable for use in official CPIs.⁴ To avoid revisions, statistical agencies can apply splicing methods, which link multilateral indexes computed over different ‘windows’ of time. While splicing eliminates the need to make retrospective updates, it reintroduces some chain drift bias. While various splicing methods have been proposed by researchers and practitioners and adopted by some NSIs, their relative performance remains debated.

We evaluate alternative splicing methods and window lengths by comparing each spliced series to its corresponding chain-drift-free multilateral index computed over all 96 months of our data. The difference between the two quantifies the chain drift bias introduced by the splicing procedure. We find that spliced multilateral indexes exhibit substantially less chain drift bias than their chained bilateral counterparts. However, the bias can still be non-negligible, making the choice of index number formula and splicing method important in practice. The GK index is particularly sensitive to the choice of splicing method.

Our results suggest that the CCDI multilateral index, calculated using a 25-month window, performs comparatively well and emerges as a leading candidate among the alternatives evaluated. The choice of splicing method has relatively little impact on the degree of chain drift bias for this index. Among the available methods, the mean splice—which averages across multiple splicing periods and is therefore less sensitive to linking on an anomalous month than alternatives—is the most attractive splicing option.

We conclude our analysis by examining the main drivers of chain drift bias in spliced multilateral indexes. We show that the most significant predictor of chain drift bias across product categories is the degree of product entry and exit, or “churn”, within the category. This analysis

⁴For instance, the U.S. Bureau of Labor Statics and the UK Office for National Statistics have policies of not revising published headline CPI figures unless a significant error is identified. This is standard practice internationally.

highlights that categories with high product turnover are where chain drift bias is likely to be most pronounced, and where longer window lengths offer the greatest potential benefits.

Our work contributes to a literature that evaluates index number methods for measuring inflation with high-frequency scanner data. This includes Feenstra and Shapiro (2003), who use scanner data for canned tuna over two years to compare the performance of fixed-base and chained bilateral indexes; Ivancic et al. (2011), who analyze 19 product categories over 15 months to quantify chain drift among superlative bilateral indexes by comparing them to a chain-drift-free GEKS-Fisher index; Melser (2018) who compares multilateral indices calculated on scanner data with different linking approaches and window lengths for eight product categories; and Diewert and Fox (2022) 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 typically been confined to a subset of available index methods and a small number of product categories over relatively short time periods. We contribute by systematically comparing the leading methods for measuring inflation using transaction data with substantially 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. This includes research quantifying the impact of product entry and exit on changes in the cost-of-living (e.g., Broda and Weinstein, 2010); measuring heterogeneity in inflation rates across households (e.g., Kaplan and Schulhofer-Wohl, 2017; Jaravel, 2019); and examining the effects of intertemporal substitution due to consumer hoarding on inflation (e.g., Ueda et al., 2024).⁵ We contribute to the strand of this literature focused on the measurement of high-frequency price dynamics, which includes work documenting

⁵Ueda et al. (2024) propose a type of price index that eliminates intertemporal substitution bias, illustrating their approach using scanner data over 30 years for processed food and daily necessities, which account for 20% of households' expenditure. They compare their method with rolling-window CCDI indexes using daily data, with windows of 7, 28 and 30 days. In contrast to our paper, they do not consider alternative multilateral indexes, splicing methods, or other potential drivers of chain drift bias.

the frequency of individual product price adjustments (e.g., Eichenbaum et al. (2011); see also the survey by Nakamura and Steinsson (2013)); and studies of high-frequency inflation during the COVID-19 pandemic (e.g., Jaravel and O’Connell, 2020).

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 evolves over time. In this section, we examine alternative approaches to doing so, with reference to the use of high-frequency data sources such as scanner or transaction data.

2.1 Bilateral index numbers

Suppose we are interested in comparing the change in the cost of the basket of goods between two sequential periods, t and $t + 1$. One way to measure 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 the construction of CPIs. If base-period quantities are used (i.e., $\mathbf{q} = \mathbf{q}^t$) the index is known as the Laspeyres index ($P_L^{t,t+1}$) and if end-period quantities are used ($\mathbf{q} = \mathbf{q}^{t+1}$) it is known as the Paasche index ($P_P^{t,t+1}$). Both indexes can be rewritten in terms of price relatives between t and $t + 1$, $\frac{p_n^{t+1}}{p_n^t}$, weighted by expenditure shares,

$s_n^t = \frac{p_n^t q_n^t}{\mathbf{p}^t \mathbf{q}^t}$. The resulting expressions are:

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

A drawback of these indexes is that they suffer from substitution bias. This arises because they use expenditure weights from only one period and therefore fail to account for the fact that when relative prices change, consumers typically substitute away from goods that have become more expensive and toward those that have become relatively cheaper.

A solution to the problem of substitution bias is offered by *superlative* indexes (see Diewert, 1976), which combine both 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 two 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 two periods— $P_W^{t,t+1}$:

$$P_F^{t,t+1} = (P_L^{t,t+1} P_P^{t,t+1})^{1/2} \quad P_{Tq}^{t,t+1} = \prod_n \left(\frac{p_n^{t+1}}{p_n^t} \right)^{0.5(s_n^t + s_n^{t+1})} \quad P_W^{t,t+1} = \frac{\sum_n \sqrt{q_n^t q_n^{t+1}} p_n^{t+1}}{\sum_n \sqrt{q_n^t q_n^{t+1}} p_n^t}.$$

2.2 Chaining and chain drift

Now suppose we are interested in tracking how the cost of the basket of goods evolves over multiple periods, and consider comparing the first to some later period $s > 2$. One approach is to construct the chain of adjacent period-to-period comparisons, $P^{1,2} \times P^{2,3} \times \dots \times P^{s-1,s}$, resulting in a *chained* index. An alternative is to make the direct, or *fixed-base*, comparison $P^{1,s}$. 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 with product churn (i.e., the entry and exit of products over time). In such cases, fixed-base comparisons rely on overlapping products between the base and

each later period, and the number of such overlapping products can decline sharply over time.⁶ For this reason, there is broad international consensus that chained indexes are preferable.⁷

A drawback of chained indexes is that, unlike fixed-base indexes, they can exhibit *chain drift*. Suppose all goods are available in each period.⁸ A chained index is said to satisfy the circularity test if $P^{1,2} \times P^{2,3} \times \dots \times P^{T-1,T} = P^{1,T}$; i.e., if the chained comparisons between the first and final period equals the direct comparison between these periods. 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. This source of bias can be particularly pronounced when indexes are constructed using transaction-level data (e.g., Ivancic et al., 2011).

Sources of chain drift. A cost-of-living index measures the expenditure a household requires to reach a fixed standard of living (Könus, 1924). The economic approach to index numbers motivates price indexes as approximations (or exact measures) of the cost-of-living index. If the price index is correctly specified to match consumers’ true preferences, the fixed-base and chained comparison between two (non-consecutive) periods will coincide; the price index will exhibit no chain drift. For all commonly used index numbers—that depend only on prices and quantities (or expenditure shares)—this equivalence requires homothetic preferences. In practice, however, the conditions required for this exact correspondence rarely hold, leading to deviations between chained and fixed-base indexes and the emergence of chain drift.

For instance, household preferences are typically not homothetic: as income rises, consumers tend to increase the share of spending on certain goods—luxuries—and reduce the share allocated to others—necessities. In such cases, the true cost-of-living index depends on compensated (i.e.,

⁶For example, the chained comparison $P^{t,t+1}$ uses data on the set of products available in periods t and $t + 1$. The fixed-base comparison $P^{1,t+1}/P^{1,t}$, instead uses data on the set of products available in both periods 1 and t (or $t + 1$). With substantial product entry and exit, the set of products available in all periods 1 and t (or $t + 1$) may comprise only a relatively small share of expenditure compared to those available in periods t and $t + 1$, particularly as t grows.

⁷For instance, the ILO (2004) CPI Manual (p. 407) notes: “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.”

⁸More precisely, suppose either all goods are available and purchased by households in each period, or that for goods that are unavailable or not purchased we observe their reservation price.

utility-constant) budget shares, which are unobserved, rather than the observed shares used in standard price indexes (Samuelson and Swamy, 1974; Theil, 1976). Preferences may also vary over time due to seasonal patterns or changes in fashion, which can cause purchases of specific goods to shift markedly from month to month. Additionally, in practice, aggregate price indexes are usually constructed using aggregate expenditure weights. However, if households have heterogeneous preferences, the resulting index may not align with the true cost-of-living index for any individual household, even under homotheticity. Another challenge is that price indexes are based on economic transactions rather than consumption, and the two do not always align. This issue is particularly pronounced in high-frequency data, due to stockpiling behavior. For example, when a product goes on sale, consumers may stock up and consume the product in later periods, reducing spending in those periods even if the price returns to its regular level.⁹ Finally, when there is product entry and exit—referred to as churn—the true cost-of-living index depends on reservation prices (i.e., the maximum price at which a consumer would still purchase the good), which are unobserved and therefore omitted from standard price indices.

Non-homotheticities, seasonal variation in preferences, aggregation across heterogeneous households, intertemporal behavior (such as stockpiling), and product churn are all potential reasons why a price index may fail the circularity test and therefore exhibit chain drift. Chain drift means the index, when chained, will fail to revert to a previous level when all prices return to the former level, and that it can drift upward or downward over time in a way disconnected from actual price changes. In high-frequency indexes, this problem can be especially severe, sometimes rendering the resulting index effectively unusable.

Dissimilarity chain linking. One approach to minimizing chain drift bias, while retaining bilateral index number comparisons, is to construct chains across periods that are most similar in terms of relative price movements. This requires a measure of “dissimilarity” of price lev-

⁹Diewert (2022) shows, for the case of a Törnqvist index, that chain drift bias arises when changes in spending shares between two periods are correlated with prices in any other period. This can occur if consumers engage in stockpiling during sales, causing future spending shares to remain low even as prices revert.

els between periods.¹⁰ Diewert (2022) recommends a “predicted share measure of relative price dissimilarity” for calculating price indexes in cases with high product churn.¹¹ Hypothetical expenditure shares are constructed using prices in period τ but quantities from period t to “predict” period t expenditure shares. For any good n , the predicted share is given by

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

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 takes values between 0 and 2. It equals 0 if prices in period τ are proportional to prices in period t (i.e., $\mathbf{p}^\tau = \lambda \mathbf{p}^t$ for some scalar λ), since then $s_{n,t} = \tilde{s}_{n,t,\tau}$ and $s_{n,\tau} = \tilde{s}_{n,\tau,t}$ for all n . This seeks to reduce chain drift by linking periods with the most similar relative prices—i.e., those closest to proportional price movements. However, as we show in Section 4, this method can still lead to substantial bias in practice.

2.3 Multilateral index numbers

Another approach to addressing chain drift bias is the use of multilateral index numbers, first proposed as a solution to chain drift bias by Ivancic et al. (2011).¹² A multilateral index computed over all periods $1, \dots, T$ is fully transitive, meaning it satisfies the circularity test. A set of multilateral indexes extend the superlative bilateral indexes defined in Section 2.1—namely, the Fisher, Törnqvist, and Walsh indexes—and like them are consistent with a flexible representation of consumer preferences and hence limit substitution bias. They are called the GEKS-Fisher index,

¹⁰This approach is also used in the context of international comparisons, where countries are linked based on similarity of their price structures. See Hill (1999, 2001).

¹¹Diewert (2009) outlines several alternative dissimilarity measures. The novelty of the predicted share measure is that it penalizes periods with limited product overlap—i.e., when many products are not observed in both periods. Diewert et al. (2022) provide empirical evidence of its potential effectiveness, particularly for seasonal products.

¹²Multilateral indexes are typically used in 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).

CCDI index and the GEKS-Walsh index respectively.¹³ The price level in period t is defined as a geometric mean of the corresponding bilateral index comparing period t with all other periods $\tau = 1, \dots, T$. Hence, the measured price level in period t under the indexes is given by:

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

The fourth multilateral index number we consider is the Geary-Khamis (GK) index,¹⁴ which differs from the other multilateral indexes in two key respects. First, it is an implicit index, defined as total expenditure divided by a volume (quantity) index. Second, it is not built upon a superlative bilateral index. Instead, it assumes a linear preference structure in which consumers regard goods as perfect substitutes. Diewert and Fox (2022) shows that the GK index is also consistent with Leontief—or perfect complements—preferences. That is, it aligns only with extreme assumptions about consumer behavior.

The GK index is defined implicitly through a set of equations that jointly determine price levels \mathbb{P}_{GK}^t , for $t = 1, \dots, T$, and a set of quality adjustment factors b^n for each good $n = 1, \dots, N$. Letting $q_n \equiv \sum_t q_n^t$ denote the total quantity of good n across all periods, the $N + T$ equations that determine the quality adjustment factors and price levels are:

$$b_n = \sum_t \left(\frac{q_n^t}{q_n} \right) \left(\frac{p_n^t}{\mathbb{P}_{GK}^t} \right) \quad \text{for } n = 1, \dots, N, \quad \mathbb{P}_{GK}^t = \frac{\mathbf{p}^{t'} \mathbf{q}^t}{\mathbf{b}' \mathbf{q}^t} \quad \text{for } t = 1, \dots, T.$$

Each adjustment factor b_n is a share-weighted average of inflation-adjusted prices for good n over all periods. The resulting price index in period t is given by total expenditure divided by the sum of quality-adjusted quantities purchased in that period.¹⁵

For each of the multilateral index numbers, it is standard to rebase the price levels relative to the first period in the data, $P_i^{1,t} = \mathbb{P}_i^t / \mathbb{P}_i^1$, for $i = GEKS - F, CCDI, GEKS - W, GK$. The

¹³The GEKS indexes are named after Gini (1931), Eltetö and Köves (1964) and Szulc (1964), while 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).

¹⁴Developed by Geary (1958) and Khamis (1970, 1972).

¹⁵The standard approach to solving this system is to iterate between the two equations until convergence. However, Diewert and Fox (2022) derives a more efficient method, based on an earlier suggestion by Diewert (1999) (see p. 360, footnote 24).

comparison of prices between any two adjacent periods, t and $t + 1$, is then given by the ratio of their respective price levels: $P_i^{t,t+1} = \mathbb{P}_i^{t+1}/\mathbb{P}_i^t = P_i^{1,t+1}/P_i^{1,t}$.

2.4 Spliced price series

Suppose a multilateral index is used to compute price levels over a fixed time period, $1, \dots, T$. When data for period $T + 1$ becomes available, recomputing the index over $1, \dots, T + 1$ will generally lead to revisions of price levels for the earlier $1, \dots, T$ periods. NSIs typically regard such revisions to past headline CPI levels as undesirable. Linking methods provide a way to address this issue.

The rolling window splice involves computing a multilateral index over an initial window of periods $t = 1, \dots, T$. When data for a new period becomes available, a new index is calculated over the updated window $t = 2, \dots, T + 1$. The price level for period $T + 1$ from this new sequence is then linked to the original series using a comparison period common to both windows (typically period T). As each new period of data becomes available, a new index is computed over the most recent T periods, and the resulting price level is spliced onto the existing series. In this approach, the window length remains fixed at T .

More concretely, suppose we compute a multilateral price series over $t = 1, \dots, T$, $\mathbb{P}_O = (\mathbb{P}_O^1, \dots, \mathbb{P}_O^T)$. For $t \leq T$, the price level is $\rho_t = \frac{\mathbb{P}_O^t}{\mathbb{P}_O^1}$. When data for period $T + 1$ becomes available, we compute a new multilateral sequence over the periods $t = 2, \dots, T + 1$, $\mathbb{P}_N = (\mathbb{P}_N^2, \dots, \mathbb{P}_N^{T+1})$. The spliced price level for period $T + 1$ is then:

$$\rho_{T+1}(\tau) = \rho_T(\tau) \times \frac{\mathbb{P}_N^{T+1}/\mathbb{P}_N^\tau}{\mathbb{P}_O^T/\mathbb{P}_O^\tau},$$

where τ is the period used to link the two index sequences. Different choices of τ correspond to alternative rolling-window splices: $\tau = T$ defines the *movement splice* (Ivancic et al., 2011); $\tau = 2$ corresponds to the *window splice* (Krsinich, 2016); and $\tau = \frac{T}{2}$ (or $\tau = \frac{T+1}{2}$ when T is odd) yields the *half splice* (de Haan, 2015)

As each subsequent period of data, $t = s + T$ (for $s > 0$), becomes available, the most recent multilateral index sequence of length T , $\mathbb{P}_{\mathcal{N}'} = (\mathbb{P}_{\mathcal{N}'}^{t-T+1}, \dots, \mathbb{P}_{\mathcal{N}'}^t)$ is appended 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 $\rho_{t-1}(\tau)$. The updated price level is given by:

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

Without additional structure on the underlying price and quantity data, there is no compelling reason for favoring any $\tau = 2, \dots, T$. Rather than selecting one period, the *mean splice*, developed by Diewert and Fox (2022), takes a geometric mean over all possible τ . This yields the normalized price level for calendar time $t > T$.¹⁶

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

A final option that we consider is to select the splicing period using a *dissimilarity* measure—such as the predicted share measure of relative price dissimilarity discussed above—by setting $\tau = \arg \min_{\tau \in 2 \dots T} \Delta_{PS}(p^T + 1, p^\tau, q^T + 1, q^\tau)$. This approach identifies the splicing period that is closest to being a proportional price change from the final period of the new window.¹⁷ This form of splicing based on relative price similarity was suggested, but not pursued, by Diewert and Fox (2022). Our paper presents the first evidence of the empirical performance of dissimilarity-based splicing for multilateral price indexes. Alternative splicing methods include fixed-base moving and expanding windows, as well as approaches that splice directly on the published series (Chessa, 2021).¹⁸ We discuss and evaluate these alternatives in the Online Appendix.

¹⁶The idea of using a mean splice was originally suggested, but not developed, by Ivancic et al. (2011), footnote 19, p. 33.

¹⁷An alternative is to select the splicing period most similar to the final period of the old window, i.e., period T . In practice, this is likely to yield very similar results when constructing monthly indexes, as in most cases months T and $T + 1$ will have similar price structures.

¹⁸Melser (2018) proposes a further alternative in which bilateral comparisons in the multilateral index number formula are weighted—assigning lower weight to comparisons between periods with less product overlap. The rolling window splice can be interpreted as a special case of this approach, with weights of one for comparisons within the window and zero for comparisons outside it.

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, the chosen linking method, and the characteristics of the underlying price and quantity data. While it is plausible *a priori* that shorter window results in greater bias, the magnitude of this effect is ultimately an empirical question. Similarly, the relative performance of different linking methods cannot be determined without empirical investigation. To address these issues, we compare a multilateral index computed over the full sample period—i.e., a fully transitive index that does not suffer from chain drift bias—with the same index computed using each of the linking procedures. The difference between these series provides a direct measure of the chain drift bias introduced by linking.

3 Scanner data

We use household scanner data from the Kantar FMCG At-Home Purchase Panel. The data capture purchases of all fast-moving consumer goods (FMCG)—including food, beverages (both alcoholic and non-alcoholic), 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 sample spans the period 2012–2019. Each year, the dataset contains purchase records from approximately 30,000 households. Households typically remain in the panel for many consecutive months. Participants record all barcoded purchases using a handheld scanner or mobile phone app. For each transaction, we observe the quantity purchased, expenditure, transaction price, and barcode-level product characteristics (including product category).¹⁹

In total, our data includes approximately 300,000 unique barcodes and over 300 million transactions, grouped into product categories. In the analysis that follows, we compute price indexes for

¹⁹The combination of transaction-level prices and expenditures, and rich product and household attributes, has made scanner data a widely used resource in economic and social science research (see Dubois et al. (2022)). See Leicester and Oldfield (2009) for a detailed description of the Kantar data, as well as comparisons with other UK data sources.

each of the 178 product categories that account for at least 0.1% of total spending over 2012–2019.²⁰ We construct monthly price indexes, treating barcodes as the elementary products in the index. We compute elementary (monthly barcode) prices by dividing the total monthly expenditure on the barcode by the total monthly quantity sold.²¹

Our data are *household* scanner data, meaning they capture transactions recorded by a sample of households. In contrast, *retail* scanner (or point-of-sale) data record all sales in a sample of stores. Both types of data enable near real-time measurement of prices and expenditures across a broad set fast-moving consumer goods. Each has distinct advantages: retail scanner data may exhibit less sampling variation if store coverage is comprehensive, while household scanner data can capture online purchases and support disaggregate inflation analysis by household type (for instance, see Chen et al. (2025)). In practice, there may be returns from combining both sources for inflation measurement.²²

Traditionally, NSIs collect price data through in-person visits by price collectors and combine this with household expenditure information from surveys such as the Consumer Expenditure Survey (CES) in the US and the Living Cost and Food Survey (LCFS) in the UK. These traditional data sources have several limitations compared to scanner data, and these limitations are reflected in standard CPI construction methods. First, price quotes are collected for only a narrow subset of products. Second, expenditure information are not available at the level of individual products,

²⁰A full list of product categories is provided in the Online Appendix. For each product category-year, we drop transactions with expenditure, volume, or unit value in the top or bottom percentile. This trimming has no substantive effect on our results.

²¹In other words, we use monthly unit values as the prices for index construction. This approach does not distinguish between variation in product availability across different regions of the UK. However, the major UK supermarkets have national store coverage and pricing policies, meaning such regional variation is likely to be small. Diewert et al. (2016) show that unit values should be calculated at the same frequency as the desired index to avoid an upward bias.

²²Both household and retail scanner data can face challenges arising from product relaunched. For example, if a product is withdrawn and reintroduced with a changed characteristic—such as a different package size—it is typically assigned a new barcode. As a result, the associated price change is not captured when using product identifiers like barcodes. This issue has been explored in the context of multilateral indexes by, e.g., Van Loon et al. (2023). A key difficulty lies in detecting relaunched when product characteristic data are limited, as is the case in our dataset. However, relaunched are generally considered more problematic for product categories such as fashion and consumer electronics than for the types of goods examined here.

but only for broad product categories—similar to the product categories in our scanner data.²³ Third, the expenditure information are available with a significant lag, often a year or more. As a result, CPIs traditionally rely on unweighted averages of a small number of price quotes to estimate product category-level prices, which are then combined using historical expenditure weights. By contrast, large scanner datasets offer several advantages, including broader coverage, more timely and disaggregate expenditure data, and the ability to compute prices at the product level—provided the issue of chain drift can be effectively addressed.²⁴

4 Comparing bilateral and multilateral indexes

One approach to calculating month-to-month price indexes using high-frequency data is to use a fixed-base Laspeyres index. Like traditional CPIs, this index uses historical spending weights.²⁵ When implemented with scanner data, however, the index can leverage detailed product-level prices and weights, covering thousands of items—an advantage over traditional approaches based on limited price quotes. Figure 4.1 plots the evolution of a fixed-base, month-to-month Laspeyres index for all fast-moving consumer goods over the period 2012-2019. The figure also plots the fixed-base superlative Törnqvist and Fisher indexes. We omit the Walsh index, as its path closely mirrors the Törnqvist index. These indexes are computed over products available in each month of the period 2012-2019.

The superlative indexes show substantially different price changes compared to the Laspeyres index, highlighting the presence of substitution bias in the latter. By the end of the first year, the Laspeyres index is around five percentage points higher than the superlative indexes—a gap

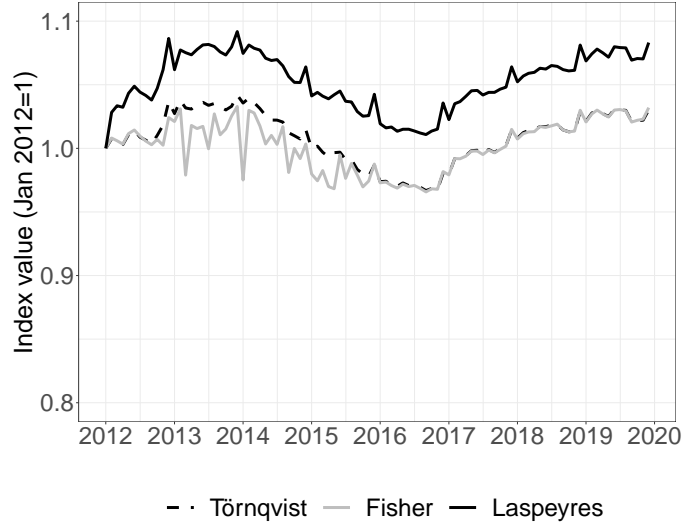
²³Implicit expenditure weights for more disaggregated items do influence index construction through sample selection, but this selection is typically updated only infrequently.

²⁴In explaining their move toward multilateral methods for CPI construction, the Australian Bureau of Statistics (2017) noted “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.”

²⁵Most NSIs use a Laspeyres-type index for CPI construction, specifically the Lowe (1823) index. In this case, the quantity weights are typically from a period prior to the price comparison, usually from the most recent expenditure survey or from an imputed update of those weights.

that persists throughout the remainder of the sample period. In the first half of data, the Fisher index displays notably more volatility than the Törnqvist and Walsh indexes. This pattern recurs in our broader results, including for the multilateral extensions of these indexes, and reflects the Fisher index’s sensitivity 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, risk becoming increasingly unrepresentative of consumer spending, as product availability evolves across seasons and over years. Figure B.1 in the Online Appendix shows the distribution of the share of spending in December 2019, across 178 product categories, that went to products also purchased in January 2012. The figure highlights significant product churn. In the median product category, only 32% of spending in December 2019 was on items with positive spending in January 2012.²⁶ Product churn on this measure is especially high in categories such as moist wipes, machine wash products, cat food, and fresh bacon joints.

Chained indexes help address product churn by requiring only that products be available in the two periods compared in each bilateral link of the chain, while allowing index weights to reflect current spending patterns. However, as discussed in Section 2.2, chained bilateral indexes

²⁶This is also equal to the weighted mean share, i.e., the share of overall spending on items across all product categories that were bought in January 2012.

constructed with high-frequency data can suffer from substantial chain drift. Figure 4.2(a) plots values for three indexes: a fixed-base Törnqvist index, a month-to-month (or period-to-period) chained Törnqvist index, and the CCDI index—the multilateral analogue of the Törnqvist index—defined in Section 2.3. We compute all three indexes using all fast-moving consumer goods. We include equivalent graphs for the Fisher and Walsh indexes, which show similar patterns, in the Online Appendix.

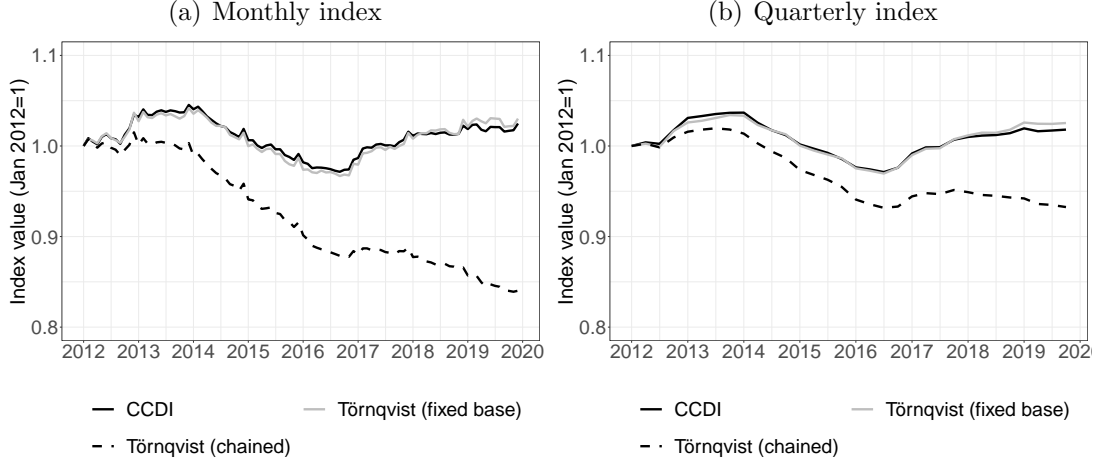
Figure 4.2(a) shows that chain drift is a significant problem for the Törnqvist index. By the end of the sample period, the chained Törnqvist index is 19 percentage points (ppt) lower than the fixed-base index, and 18 ppt lower than the multilateral CCDI index. In contrast, the CCDI index is much closer in value to fixed-base index than the chained Törnqvist index.

Panel (b) of Figure 4.2 replicates this comparison using quarterly indexes. The quarterly CCDI and fixed-base Törnqvist index resemble their monthly counterparts, though they smooth over short-term month-to-month variation. The divergence between the chained Törnqvist and the other two indexes is smaller at the quarterly frequency, but the chained index still records considerably lower cumulative inflation.

The monthly chained Törnqvist index for all items falls by 16 ppt over the full period—an implausibly large decline. Figure B.3 in the Online Appendix shows the distribution of cumulative price changes across all 178 product categories. The chain drift problem is even more pronounced: ten categories record price changes of less than -50%, while nine record increases of more than 20%.

Compared with the fixed-base Törnqvist index, the CCDI index has the advantage of not requiring products to be present in a single base period. As a result, it remains more representative of consumer spending over time. Nonetheless, Figure 4.2 shows that over 2012–2019 the CCDI and fixed-base Törnqvist indexes yield a similar picture of inflation in fast-moving consumer good inflation. However, in general, this need not be the case.

Figure 4.2: *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.

Panel (a) of Figure 4.3 compares the distribution of differences between the final-period values of the CCDI index and three alternative Törnqvist indexes: a period-on-period chained index, a chained index using the predicted share dissimilarity method, and a fixed-base index. Each distribution is shown as a box plot, where the box represents the interquartile range, the line marks the median, and the whiskers extend to 1.5 times the interquartile range below and above the box. Outliers beyond this range are plotted individually.

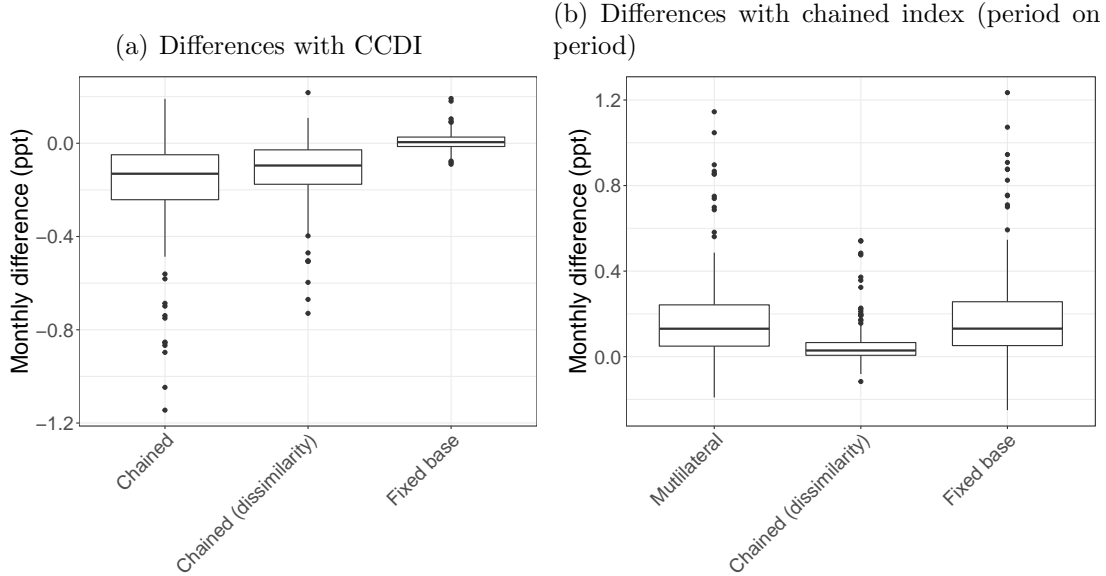
Comparing the CCDI with the chained Törnqvist index suggests that chain drift is negative for most products. When the Törnqvist index is chained using the standard period-on-period approach, 22 out of 178 product categories exhibit positive chain drift bias.²⁷ Differences between the CCDI and the Törnqvist index chained using the dissimilarity index are smaller, but still substantial. The median average monthly difference for the dissimilarity-chained Törnqvist index is -0.09 ppt, compared to -0.13 ppt for the period-on-period chained Törnqvist index.

The differences between the CCDI and the fixed-base Törnqvist index are smaller than those for the chained indexes, but in some cases, they remain substantial. These cases typically involve

²⁷Notably, fresh fruits account for three of the four product categories with the highest positive chain drift bias: citrus fruits, apples, chilled breads, and pears.

product categories with a high degree of product churn, reflecting bias introduced as the fixed-base index becomes increasingly unrepresentative over time.

Figure 4.3: *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 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 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.

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 4.3 shows the differences between the period-on-period chained Törnqvist index and the alternative approaches. It shows that the dissimilarity-based chaining yields results that are generally very similar to period-on-period chaining, which, as discussed, tends to imply implausible index changes.

Closer inspection of the dissimilarity-based chaining approach reveals that, in many cases, the most similar period—according to the dissimilarity measure—is the immediately preceding one. As a result, this method often yields outcomes similar to standard period-on-period chaining. To illustrate this, Figure B.4 in the Online Appendix plots the time paths of the different indexes for a specific product category: Chocolate and Confectionery. This category exhibits particularly high

chain drift bias—recording the second-largest difference between the CCDI index (or the fixed-base index) and the chained Törnqvist—and also shows one of the largest reductions in chain drift bias when using the dissimilarity approach relative to the period-on-period method. Despite this, the reduction remains modest. From 2012 to 2016, the dissimilarity-chained index closely tracks the period-on-period Törnqvist index, as both rely on the same chaining periods. It is only from 2016 onward, when a different reference period is selected, that the two indexes begin to diverge. Figure B.4 also shows that while the fixed-base index yields a final value similar to the CCDI, its trajectory over time is substantially more volatile—reflecting the high degree of seasonal product churn in this category.

4.1 Different multilateral indexes

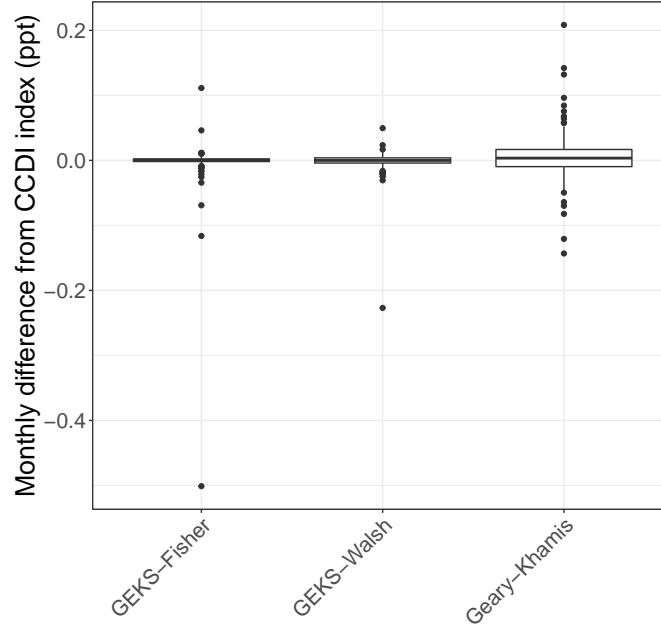
In this section, we quantify the difference in measured inflation across the four multilateral price indexes discussed in Section 2.3—the CCDI, GEKS-Fisher, GEKS-Walsh and GK indexes. For each index and product category, we compute a price index spanning the full sample period, i.e., all 96 year-months of data. Because this calculation does not entail any linking, the resulting indexes are free from chain drift. To summarize inflation differences, we calculate, for each of the GEKS-Fisher, GEKS-Walsh, and GK indexes, the difference in average monthly price changes (up to December 2019) relative to the CCDI index.²⁸ These differences form three distributions—one for each index comparison—which we present as box plots in Figure 4.4.

The differences between the GEKS-Fisher and GEKS-Walsh indexes and the CCDI index in the final period are small for the vast majority of product categories, which is perhaps unsurprising given that the bilateral forms of these indexes approximate each other to the second order. However, for a few categories, the difference is substantial. For example, the GEKS-Fisher index reports an average monthly inflation rate 0.5 ppt lower than the CCDI index for chilled flavored milk, and 0.11 ppt lower for ambient cakes and pastries. When cumulated over the full 96-months,

²⁸The average monthly inflation rate for a given index over the 96 months is calculated as $x^{\frac{1}{96}} - 1$ where x is the final period value of the index.

these gaps translate into final period index differences of 31 ppt and 13 ppt, respectively. The GEKS-Walsh index also yields lower inflation rates for these items than the CCDI, with differences of 0.22 ppt and 0.02 ppt, respectively. Conversely, the GEKS-Fisher index records an average monthly inflation rate that is 0.11 ppt higher than the CCDI index for other vegetables, while the GEKS-Walsh shows a smaller but still positive difference of 0.02 ppt for the same category.

Figure 4.4: *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.

These occasional differences between the CCDI and the other GEKS indexes appear to stem from anomalously large price and quantity changes for specific items occurring in a single month, which then have a persistent effect on the cumulative index. In such cases, because expenditure shares change less than the underlying quantities and prices, the CCDI is less affected than the quantity-based GEKS-Fisher and GEKS-Walsh indexes. This suggests that the CCDI may be less sensitive to certain extreme values.²⁹

²⁹For example, in the case of chilled flavored milk, the CCDI index falls by 2.5% in a single month, while the GEKS-Walsh index drops by 7%, and the GEKS-Fisher index by an implausible 43%. This accounts for most of the difference in their final index values. Closer inspection indicates that this was due to a sharp change in recorded quantities for three products—possibly the result of a change in units of measurement—while expenditure levels

The differences in average monthly inflation rates between the GK and CCDI indexes are generally larger than those observed between the CCDI and either the GEKS-Fisher or GEKS-Walsh indexes. The interquartile range of differences for the GK index spans from -0.01 ppt to 0.02 ppt, compared to narrower ranges of -0.002 to 0.002 ppt for the GEKS-Fisher and -0.004 to 0.003 for the GEKS-Walsh. The GK index also registers a significantly *higher* price increase for chilled flavoured milk relative to the CCDI index—by contrast to the GEKS indexes, which record substantially lower inflation for this category. Specifically, the GK index reports an average monthly inflation rate that is 0.2 ppt higher than the CCDI index.

Our analysis suggests that the CCDI index compares favorably to the other multilateral indexes considered for measuring inflation using high-frequency data. The most pronounced advantage is over the GK index: while the CCDI and other GEKS-type indexes tend to produce similar results, the GK index often diverges. This is perhaps unsurprising, as the GK index is not based on a superlative bilateral index and instead is consistent with unrealistic assumptions about consumer preferences. The case for preferring the CCDI index over the other GEKS-type indexes is more nuanced, but our findings indicate that the CCDI index is less sensitive to large outliers. NSIs might therefore prefer the CCDI as it is more robust to anomalous price changes that could be measurement errors.

In practice, however, NSIs cannot use multilateral indexes calculated over the full sample periods for headline CPIs, as doing so would require revising historical values whenever new data become available. Hence, we next examine rolling-window implementations of these indexes, which introduce an additional layer of complexity: outcomes now depend not only on the index number formula but also on the linking method used to splice together successive windows.

remained relatively stable. As a share-weighted index, the CCDI was much less sensitive to these changes than the quantity-based GEKS-Fisher and GEKS-Walsh indexes.

5 Chain drift bias in spliced indexes

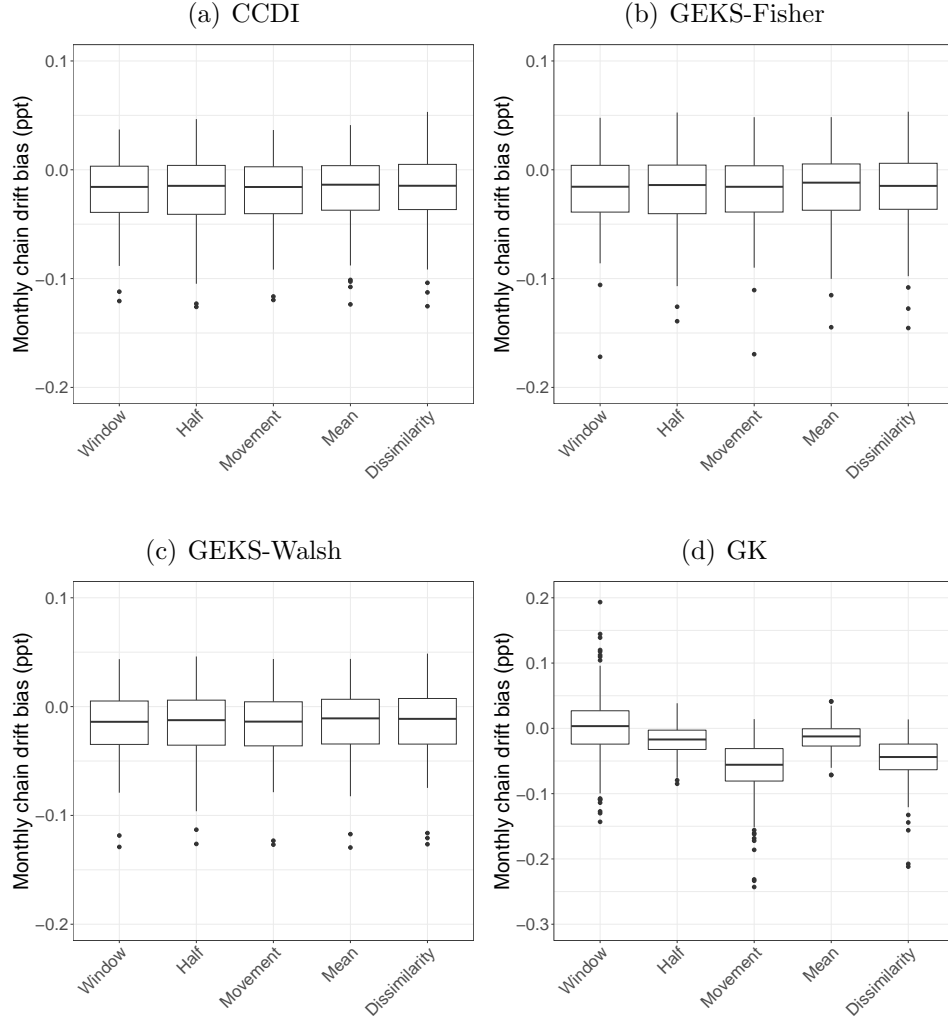
In this section, we quantify the chain drift bias resulting from different linking procedures used to extend multilateral indexes, as discussed in Section 2.4. We also examine how chain drift bias varies with different window lengths used in rolling-windows methods. To do this, we compare the average monthly inflation rates from spliced indexes for each multilateral index number with those from the non-spliced series, which is computed using all 96 year-months of data. Since the non-spliced series is fully transitive, it is free from chain drift, providing a direct measure of chain drift bias in the spliced series. We undertake this comparison for each product category across the CCDI, GEKS-Fisher, GEKS-Walsh, and GK index numbers.

5.1 Splicing methods

We first fix the window length at 25 months and compare different linking methods: the window, half, movement, and mean splices. Additionally, we use the predicted share measure of relative price dissimilarity to select the splicing period. Further results, including those 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.

Figure 5.1 presents boxplots summarizing the distribution of differences between spliced indexes and the non-spliced index (which uses all periods of data) across product categories. Each of the four panels corresponds to a different multilateral index number, and each boxplot within a panel represents a different splicing method. Some product categories exhibit a very high degree of chain drift bias. To make the plots easier to read, we exclude these extreme cases by removing the three product categories with the largest positive values of chain drift, as well as the three with the largest negative values, from 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 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 showing positive bias. The median bias ranges from -0.02 ppt with the movement splice to -0.01 ppt with the mean and half splices, while the interquartile range of biases spans 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 similar results for the GEKS-Walsh and GEKS-Fisher indexes, with chain drift bias patterns closely resembling those for the CCDI index. However, extreme cases of chain drift bias are more common for the GEKS-Fisher index.

Panel (d) shows that the GK index is more sensitive to the choice of 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 monthly chain drift bias for the movement splice is -0.06 ppt, much greater than for the other splicing methods. The window splice also leads to greater dispersion in chain drift bias under the GK index compared to other indexes. In addition, the most extreme outliers (not shown in the plots) are much higher for the GK index. For instance, 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.³⁰

Figure 5.1 provides the first empirical evidence on the use of the predicted share dissimilarity method for splicing rolling-window multilateral indexes. In the first three panels, its performance is very similar to the other methods. These results are more encouraging than for its use in chaining bilateral indexes; see Figure 4.3 in Section 4.

Overall, our findings suggest that the CCDI and GEKS indexes are robust to the choice of extension method, as the distribution of chain drift bias is similar across these methods. In contrast, the GK index is more sensitive to the extension method, with the half and mean splice yielding the best results. The mean splice avoids the risk of linking solely on a period that may happen to exhibit unusual spending and price patterns, a potential issue with the window, half and movement splicing methods. The predicted share dissimilarity method, however, is less user-friendly because the linking period is not known *ex ante* and varies between pairs of windows being linked, making it harder to explain to users. Considering both the empirical evidence in Figure 5.1 and these practical considerations, the mean splice seems to be the preferred choice.

³⁰The index for seasonal biscuits also exhibits some chain drift when calculated using the CCDI index, but the bias is much smaller: -0.05 ppt with the window splice and -0.06 ppt with the movement splice. Using the mean or half splice mitigates the extreme chain drift bias for this product category with the GK index: the biases are -0.28 ppt with the mean splice and -0.16 ppt with the half splice.

5.2 Drivers of chain drift bias

The results from the previous section emphasize the importance of using relatively long window lengths for spliced multilateral index numbers to minimize chain drift bias. In the section, we assess the circumstances under which chain drift bias is likely to be a significant issue. Specifically, we consider how chain drift bias at different window lengths relates to five potential drivers, each measured separately for each product category. *Monthly churn*: This is the share of spending on products in the current month that were not observed being purchased in the previous month.³¹ *Annual churn*: This is 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 potential cause of chain drift (Melser and Webster, 2021). *Seasonality in pricing* (*‘weak seasonality’*): We measure this by estimating a regression of log price on product fixed effects and month dummies for each product category. We measure the degree of seasonality as the difference between the largest and smallest month dummy coefficients. *The frequency of price promotions*: This refers to the percentage of transactions each year that involve price promotions. *The frequency of quantity promotions*: This refers to the percentage of transactions each year that involve quantity promotions, such as two-for-one offers.

Table B.6 in the Online Appendix shows the distribution of these measures across product categories. Monthly churn is highest for the category seasonal biscuits, with an average of 14.6% of spending each month directed toward products not purchased in the previous month. This category also exhibits the highest annual churn and seasonal pricing. Annual churn is notably high for products like chocolate and air fresheners, while monthly churn is more prominent for seasonal categories like vitamins, minerals and skincare. Apart from seasonal biscuits, seasonal pricing is significant for soft fruits and fortified wines. Price promotions are most prevalent in ‘mini

³¹Multilateral methods such as the GEKS-type indexes rely on there being some degree of product overlap for every pair of months considered. A lack of matching can cause the resulting indexes to be unreliable. Diewert and Shimizu (2023) show that product churn can be a particular problem with some high-tech products, such as laptops.

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 through the following analysis. For each multilateral index number, we calculate the absolute value of the cumulative chain drift bias over all 96 months for a spliced series (using the mean splice) computed with a 25-month rolling window. We then regress this measure on each of the potential drivers of chain drift bias described above.³² Each observation in the regression corresponds to a product category (excluding the three categories with the highest levels of chain drift bias). We also report results using a 7-month window in Table B.7 in the Online Appendix. These results suggest that pricing seasonality plays a more important role in driving chain drift bias when shorter window lengths are used—likely because such windows that do not span at least 13 months do not allow like-for-like comparisons of prices at the same point in the seasonal cycle.

Table 5.1 shows that the effects of seasonal pricing and promotions are small and statistically insignificant. Instead, higher rates of product churn emerge as the main determinant of chain drift bias. For the CCDI, GEKS-Fisher, and GEKS-Walsh indexes, each percentage point increase in annual churn is associated with a 0.15 to 0.17 percentage point increase in chain drift bias—a sizeable effect, given the relatively low overall bias at a 25-month window length. The GK index stands out as an exception. Unlike the other indexes, it is more sensitive to monthly churn than to annual churn. Each percentage point increase in annual churn is associated with a 0.06 ppt increase in chain drift bias for the GK, while each percentage point increase in monthly churn is associated with a 0.11 ppt increase—although the estimates for monthly churn are imprecise. The results suggest that product churn is a key determinant of chain drift bias. For the CCDI index spliced using a 25-month window, monthly and annual churn together account for 88% of the total explained variance in chain drift bias. Longer window lengths mitigate the impact of

³²We weight the regressions by the number of goods in each product category, to account for sampling uncertainty in the right-hand-side variables and potential heteroskedasticity. Results estimated using unweighted OLS are very similar.

annual churn and seasonal prices on the CCDI and GEKS indexes. In contrast, high-frequency churn (at the monthly level) poses a particular challenge for the GK index.

Table 5.1: *Determinants of chain drift bias (25 month window length)*

	CCDI	GEKS-Fisher	GEKS-Walsh	GK
Monthly churn	0.050 (0.130)	0.055 (0.142)	−0.012 (0.124)	0.116 (0.084)
Annual churn	0.166*** (0.051)	0.171*** (0.055)	0.149*** (0.048)	0.068** (0.034)
Pricing seasonality	0.010 (0.050)	0.028 (0.055)	0.008 (0.048)	−0.038 (0.033)
Price promotions	−0.0004 (0.028)	−0.012 (0.031)	0.003 (0.027)	−0.024 (0.018)
Quantity promotions	−0.029 (0.019)	−0.041** (0.021)	−0.035* (0.018)	−0.003 (0.012)
Observations	175	175	175	175
R ²	0.127	0.122	0.101	0.072

*Note: All indexes are extended using the mean splice. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$. The regressions are weighted using the square root of the average number of observations in each calendar month for each product category.*

6 Conclusions

Accurate and timely inflation measurement is essential for effective design and implementation of economic policy—particularly in times of high and volatile inflation. The growing availability of transaction-level data offers the potential to construct inflation measures at high frequencies, enabling faster collection, compilation, and publication of CPIs, and drawing on a richer set of underlying information than traditional methods. A range of approaches for utilizing such data have been proposed—and, in some cases, adopted—by NSIs. However, to date there is a lack of systematic empirical evidence comparing these methods. Such comparisons are crucial, as different index number formulae and extension methods can yield materially different results.

In this paper, we help fill this evidence gap by providing systematic, quantitative comparisons of competing methods for measuring month-to-month inflation using long-run transaction-level scanner data. Our analysis covers 178 product categories within the fast-moving consumer good segment of the economy. We include in our comparisons a number of proposed methods that, until now, have seem limited or no empirical applications.

We find that, among the methods examined, the CCDI multilateral index—implemented with a 25-month rolling window and mean splice linking—performs particularly well. In addition, we provide novel evidence on the determinants of chain drift bias in spliced multilateral indexes. Our results highlight product churn as a key driver of chain drift bias.

There are several potential avenues for future research. One important question concerns the variation in chain drift bias across months within a given index series. While our analysis focuses on the average monthly chain drift bias over a long time period, the month-to-month variability of the bias, and how this varies across low and high inflation environments, also affects the reliability of multilateral indexes for real-time inflation measurement. Future work could also explore the role of disappearing products in shaping index behavior, as well as the effectiveness of imputation methods for handling 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. *Economic Journal* 92(365), 73–86.
- Chen, T., P. Levell, and M. O’Connell (2025). Measuring cost of living inequality during an inflation surge. *CEPR Discussion Paper* (DP19388).
- 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 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* 161, 36–46.
- Diewert, W. (1976). Exact and superlative index numbers. *Journal of Econometrics* 4(2), 115–145.
- Diewert, W. (1999). Axiomatic and economic approaches to international comparisons. In A. Heston and R. Lipsey (Eds.), *International and Interarea Comparisons of Income, Output and Prices*, Volume 61 of *Studies in Income and Wealth*, pp. 13–87. Chicago: The University of Chicago Press.
- Diewert, W. (2009). Similarity indexes and criteria for spatial linking. In D. P. Rao (Ed.), *Purchasing Power Parities of Currencies: Recent Advances in Methods and Applications*, pp. 183–216. Cheltenham, UK: Edward Elgar.
- 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.
- Diewert, W. E. and C. Shimizu (2023). Scanner Data, Product Churn and Quality Adjustment. Vancouver School of Economics, University of British Columbia Discussion Paper 23-06.
- Dubois, P., R. Griffith, and M. O’Connell (2022). The Use of Scanner Data for Economics Research. *Annual Review of Economics* 14(1), 723–745.
- 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.
- Feenstra, R. C. and M. D. Shapiro (2003). High-Frequency Substitution and the Measurement of Price Indexes. In *Scanner Data and Price Indexes*, pp. 123–149. Chicago: University of Chicago Press.
- 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.
- Gini, C. (1931). On the circular test of index numbers. *Metron* 9(9), 3–24.
- Hill, R. J. (1999). 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.
- Könus, A. A. (1924). The Problem of the True Index of the Cost of 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.
- Leicester, A. and Z. Oldfield (2009). Using scanner technology to collect expenditure data. *Fiscal Studies* 30(3/4), 309–337.
- Lowe, J. (1823). *The Present State of England in Regard to Agriculture, Trade and Finance*, second ed. London: Longman, Hurst, Rees, Orme and Brown.
- Melser, D. (2018). Scanner data price indexes: Addressing some unresolved issues. *Journal of Business & Economic Statistics* 36(3), 516–522.
- 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.
- Samuelson, P. A. and S. Swamy (1974). Invariant Economic Index Numbers and Canonical Duality: Survey and Synthesis. *The American Economic Review 64*(4), 566–593.
- Szulc, B. (1964). Indices for Multiregional Comparisons. *Przegląd Statystyczny 3*, 239–254.
- Theil, H. (1976). *Theory and Measurement of Consumer Demand*, Volume 2. Amsterdam: North Holland.
- Ueda, K., K. Watanabe, and T. Watanabe (2024). Household Inventory, Temporary Sales, Price Indices. *International Economic Review 65*(1), 217–251.
- Van Loon, K., A. Mierop, and D. Roels (2023). Using Multilateral Hedonic Methods to Capture Product Relaunches. Meeting of the Group of Experts on Consumer Price Indices, 7-9 June.

APPENDIX: FOR ONLINE PUBLICATION

Inflation measurement with high-frequency data

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

July, 2025

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 FMBW, the spliced price level in period t is given by:

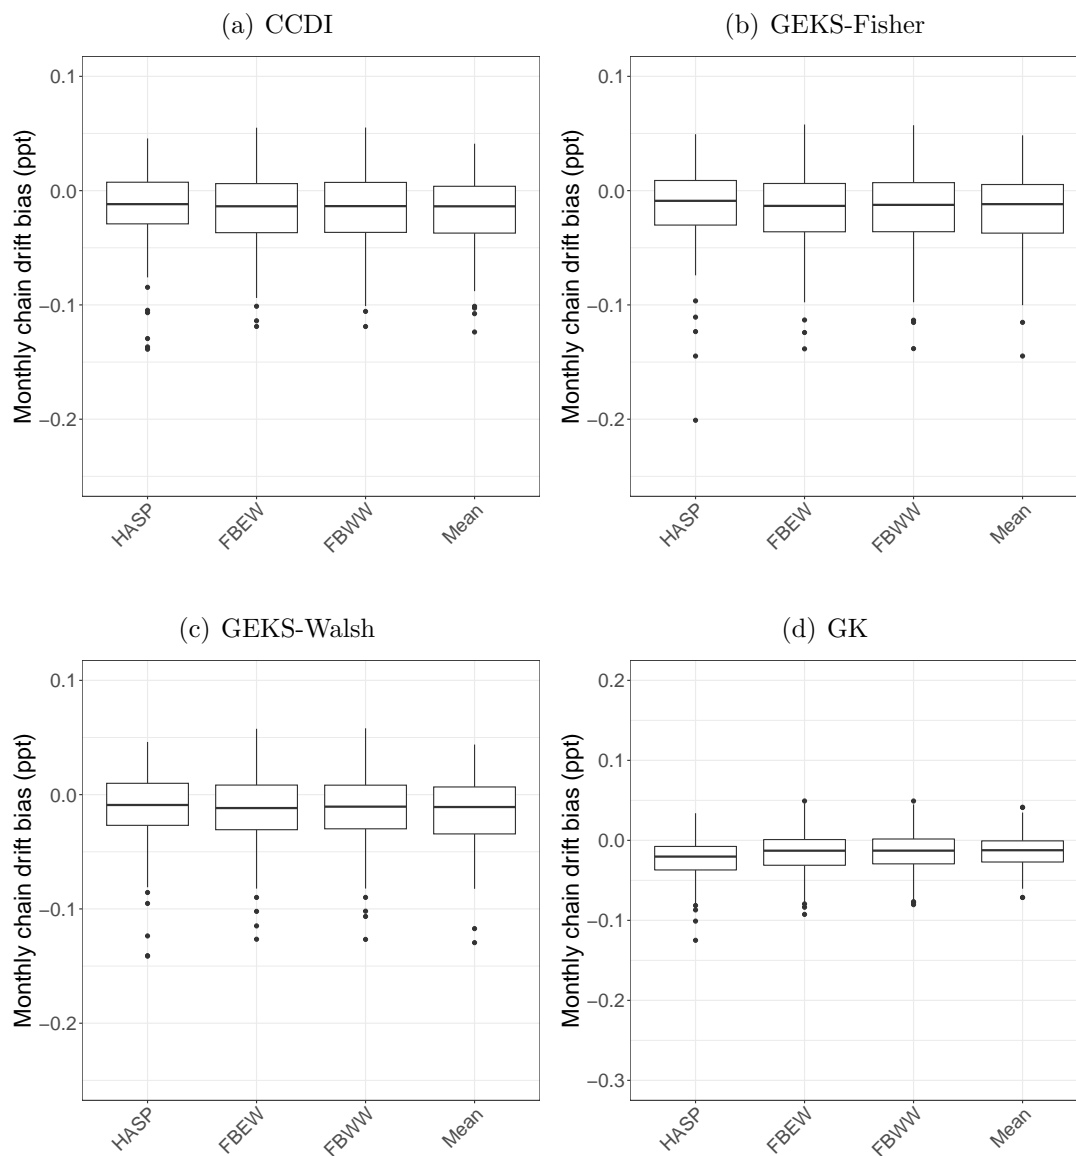
$$\rho_t = \rho_t \frac{\mathbb{P}_t^t}{\mathbb{P}_t^t},$$

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 A.1 shown plots the distribution of monthly chain drift biases associated with the half-splice on the published series, FBEW and FBMW approaches.

Figure A.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.

B Additional Figures and tables

Table B.1: *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 Frais	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

Table B.2: *Product categories (2)*

	Annual spending share (%)		
	mean	min	max
Soft Fruit	2.25	1.88	2.55
Tropical Fruits	0.49	0.40	0.55
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

Table B.3: *Product categories (3)*

	Annual spending share (%)		
	mean	min	max
Frozen Vegetarian Prods	0.22	0.19	0.26
Other Frozen Foods	0.16	0.15	0.18
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

Table B.4: *Product categories (4)*

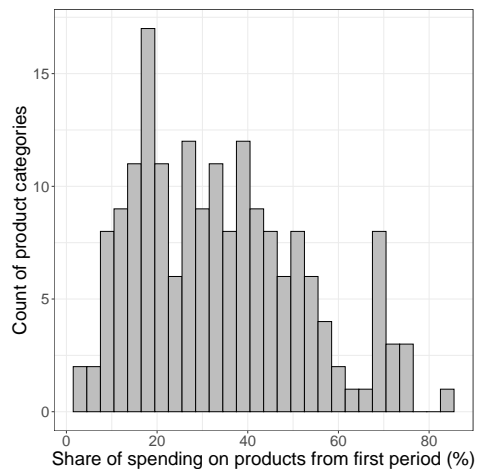
	Annual spending share (%)		
	mean	min	max
Everyday Treats	0.40	0.37	0.42
Healthier Biscuits	0.24	0.21	0.25
Savoury Biscuits	0.15	0.13	0.18
Seasonal Biscuits	0.13	0.12	0.14
Special Treats	0.16	0.14	0.17
Frozen Confectionery	0.35	0.32	0.40
Total Ice Cream	1.03	0.94	1.16
Chocolate Confectionery	2.47	2.38	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

Table B.5: *Product categories (5)*

	Annual spending share (%)		
	mean	min	max
Air Fresheners	0.36	0.33	0.39
Batteries	0.25	0.24	0.26
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

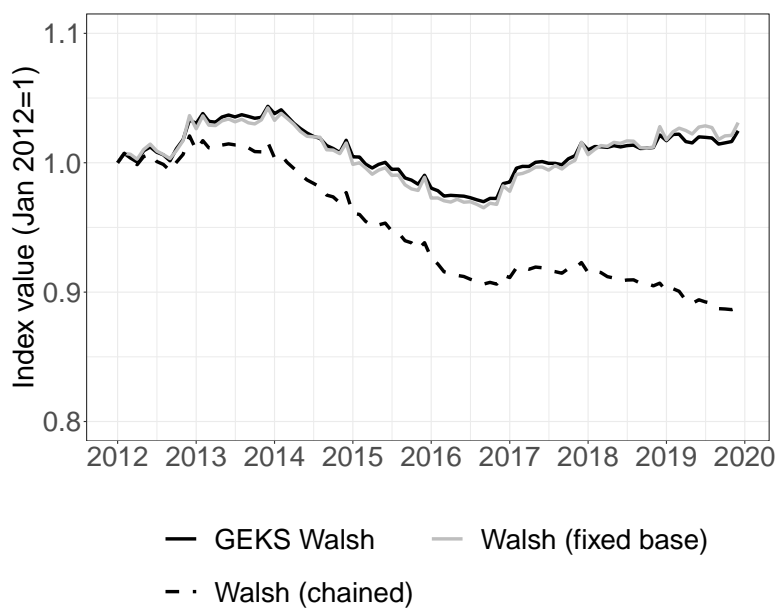
Note: Authors' calculations based on the Kantar FMCG At-Home Purchase Panel for 2012-2019.

Figure B.1: *Share of final period spending on products available in first period*



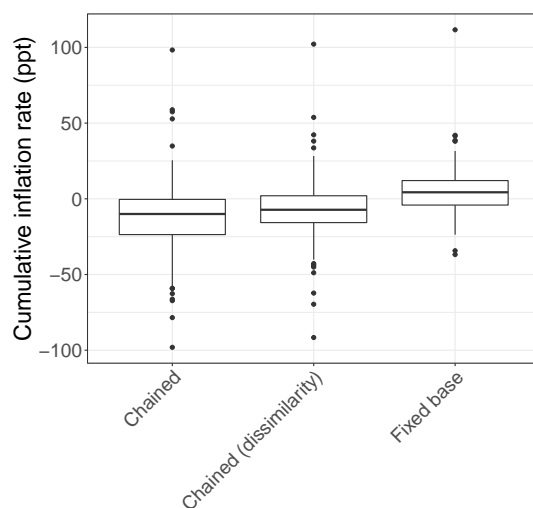
Note: Figure shows the distribution of the share of spending in the final period (December 2019) that goes on products that were purchased in the first period (January 2012) across product categories.

Figure B.2: *Chain drift bias: GEKS Walsh vs Bilateral Walsh*



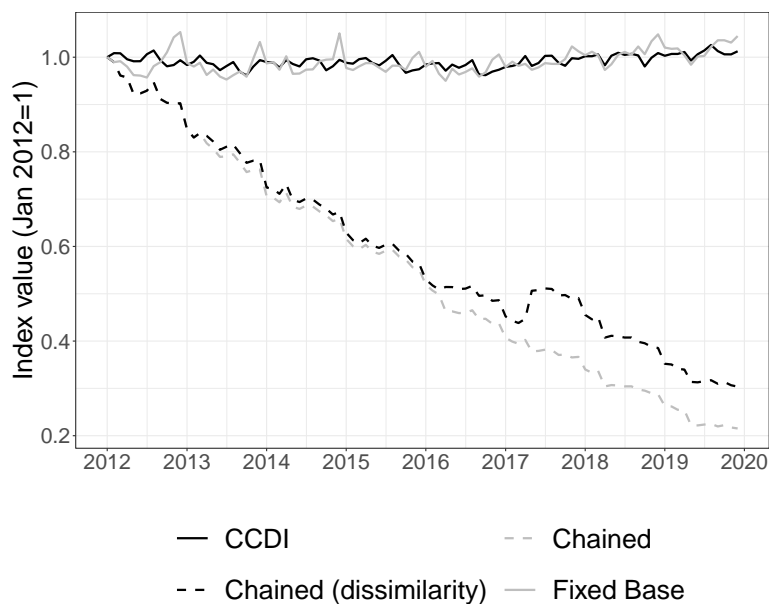
Note: Figure shows index number values for the GEKS-Walsh multilateral index, the Walsh fixed base index and a monthly chained Walsh index. The indexes are calculated across all fast-moving consumer goods.

Figure B.3: *Distribution of cumulative inflation rates for the Törnqvist index with different chaining methods*



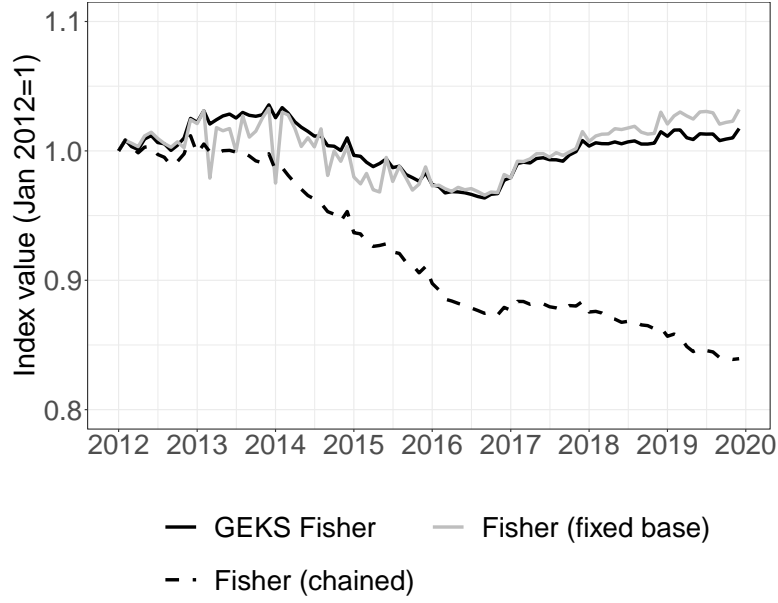
Note: Figure shows the distribution of cumulative price changes from January 2012 to December 2019 across 178 product categories calculated using the bilateral Törnqvist index using different monthly chaining methods.

Figure B.4: *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.

Figure B.5: *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.

Table B.6: *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.

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

	CCDI	GEKS-Fisher	GEKS-Walsh	GK
Monthly churn	−0.006 (0.108)	0.019 (0.115)	0.055 (0.091)	−0.003 (0.098)
Annual churn	0.356 (0.275)	0.323 (0.294)	0.334 (0.233)	0.560** (0.249)
Pricing seasonality	0.504*** (0.109)	0.503*** (0.116)	0.421*** (0.092)	0.402*** (0.099)
Price promotions	−0.058 (0.058)	−0.054 (0.062)	−0.041 (0.049)	−0.097* (0.053)
Quantity promotions	0.042 (0.039)	0.044 (0.042)	0.020 (0.033)	0.036 (0.035)
Observations	175	175	175	175
R ²	0.206	0.183	0.208	0.215

*Note: All indexes are extended using the mean splice. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$. The regressions are weighted using the square root of the average number of observations in each calendar month for each product category.*

C Different window lengths

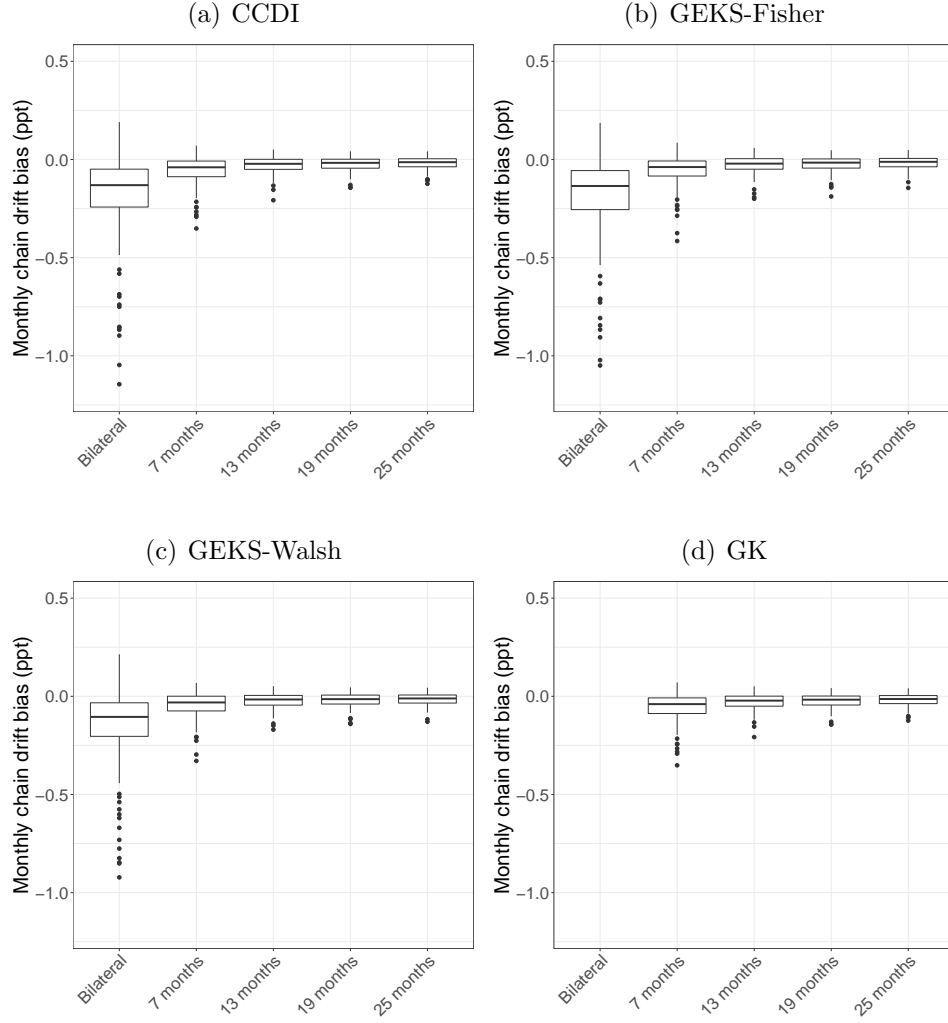
In Figure C.1 we summarize the impact of different window lengths on average monthly chain drift bias, in all cases using the mean splice. The figure is structured similarly to Figure 5.1 in the main text – each panel represents a different index number, and within each panel, the boxplots correspond 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 C.1 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 C.1: *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 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.