

Article

Product grouping: measuring inflation in dynamic clothing markets

Clothing indices can be downwardly biased because of products rapidly entering and exiting the market. Traditional methods that account for these biases cannot be replicated with our large, web-scraped datasets. This article explores a new method - product grouping.

Contact:
Helen Sands
cpi@ons.gov.uk
+44 1633 456900

Release date:
9 November 2021

Next release:
To be announced

Table of contents

1. [Main points](#)
2. [Web-scraped clothing data](#)
3. [Challenges associated with measuring clothing inflation](#)
4. [Product grouping](#)
5. [Assessing product grouping](#)
6. [Attribute-based product grouping](#)
7. [Results](#)
8. [Future developments](#)
9. [Related links](#)

1 . Main points

- In this article we explore using web-scraped data for measuring clothing prices to increase our coverage of online-only retailers, coverage of products per retailer, and frequency of collection.
- The pace that products enter and exit clothing markets makes it challenging to appropriately measure their changing prices.
- To account for these challenges, we are researching “product grouping” for use with web-scraped data, where we track the average price of groups of similar products over time.
- Preliminary results suggest that product grouping, when done appropriately, may reduce downward biases introduced into price indices for clothing.
- While alternative data sources are being [introduced into consumer price statistics from 2023](#), these ongoing research and developments for clothing using alternative data sources mean that we will not be implementing these sources for clothing until 2024, at the earliest.

2 . Web-scraped clothing data

Alternative data sources, namely scanner and web-scraped data, and methods to use these data sources are being [introduced into the production of UK consumer price statistics](#) from 2023. Traditionally, our clothing prices are collected from physical outlets, but increased spending in online-only clothing retailers in recent years means our market share coverage has been falling.

Web-scraped clothing data provide several advantages to traditionally collected price data, including increased coverage of online-only retailers, but also increased product coverage per retailer and more frequently collected data. In any one month our web-scraped clothing data contain approximately 500,000 unique products. This compares with approximately 20,000 clothing products collected each month using traditional methods.

In future, web-scraped clothing data will be integrated with clothing data collected using traditional methods to provide a broad market coverage of online retailers, high street chains and smaller or independent retailers. Details of how price indices constructed from different data sources will be combined is discussed further in [Introducing alternative data into consumer price statistics: aggregation and weights](#).

Web-scraped data do not contain product-level weights. We are exploring [methods to approximate weights](#). However, this is ongoing research and results in this paper remain unweighted.

3 . Challenges associated with measuring clothing inflation

The nature of clothing markets results in measurement challenges primarily due to the pace that products enter and exit the market. [These challenges have also been experienced historically \(PDF, 61KB\)](#) but, with the proposed move to include web-scraped prices, we must develop new methods to cope with the larger scale and complexity of these data.

[Inflation is traditionally measured](#) by tracking prices of individual products and aggregating them through time. Clothing products rarely exist for more than a few months, they experience high "product churn". In our web-scraped clothing data, around 150,000 of the 500,000 (approximate) products leave and enter the market each month.

Table 1 shows how, under a simple example of these conditions, price indices can be unrepresentative as they only use products that exist in both periods.

Table 1. The August price index only uses floral winter dress 3 since it is the only product that has prices in both January and August

Product	Price, Jan	Price, Aug	Price change (Aug/Jan)
Floral winter dress 1	18		Cannot form
Floral winter dress 2	18		Cannot form
Floral winter dress 3	24	18	0.75
Floral summer dress 1	60		Cannot form
Floral summer dress 2		45	Cannot form
Floral summer dress 3		45	Cannot form
Party midi dress 1	100		Cannot form
Party midi dress 2		90	Cannot form
			Price index 0.75

Source: Office for National Statistics


Furthermore, clothing products typically enter the market at a high price and leave at a low price, often on a clearance sale. This creates an implicit price increase when products are replaced which is not normally captured by index methods (as shown in Figure 1). In our web-scraped data, prices of individual products fall on average by 3% between two consecutive months.

Figure 1. Index methods may only capture falls in prices of individual product lines (shaded lines) and not price resetting when product lines are replaced (dotted lines)

Figure 2 shows a GEKS-Jevons index ([our currently preferred unweighted method](#)) on web-scraped women's dresses data, without applying methods to account for these measurement challenges. Within just seven months, prices fall 20%: the index captures the declining price of individual products, but not implicit price resetting that occurs when new equivalent product lines are released.

Figure 2: An unadjusted price index for women's dresses using web-scraped data falls by 20% in seven months

Figure 2: An unadjusted price index for women's dresses using web-scraped data falls by 20% in seven months



Source: Office for National Statistics

Traditionally, we account for clothing measurement challenges by manually linking products leaving the market with suitable replacement products. However, it is unfeasible to scale this manual replacement process for the 150,000 (on average) leaving products every month when using web-scraped data. Instead, we explore using a method we refer to as product grouping.

4 . Product grouping

Web-scraped data are first [classified to clothing consumption segments](#), such as women's dresses. Within consumption segments, groups are formed of similar products, such as floral summer dresses. We then track the average price of each group instead of each individual product. We have explored similar concepts previously in [Clustering large datasets into price indices](#). Countries such as [the Netherlands](#) (PDF, 860KB) and [Belgium](#) (PDF, 436KB) have also explored similar applications.

Table 2 shows a stylised example of how we could group a sample of dresses, based on their product names, to reduce the effect of product churn.

Table 2. Product grouping measures changes in average prices of product groups, allowing more prices to contribute to inflation

Product	Price, Jan	Price, Aug	
Floral winter dress 1	18		
Floral winter dress 2	18		
Floral winter dress 3	24	18	
Floral summer dress 1	60		
Floral summer dress 2		45	
Floral summer dress 3		45	
Party midi dress 1	100		
Party midi dress 2		90	
			Price change (Aug/Jan)
Average price: Floral winter dresses	20	18	0.9
Average price: Floral summer dresses	60	45	0.75
Average price: Party midi dresses	100	90	0.9
			Price index ¹ 0.85

Source: Office for National Statistics

Notes

1. These groups have been given an equal weight within this index but in future we may look to weight groups by the number of products they contain or by another suitable approximate.

Since price changes are measured within product groups, that are more likely to be available throughout the year than individual products, product churn is reduced. Furthermore, since products entering and leaving the market influence the group average price, the implicit price resetting as new products enter the market is also captured.

However, using product grouping introduces a new measurement challenge. In Table 3, all floral dresses are combined into a single group. This combined group shows a price increase, however, when considered independently, floral summer and winter dresses both show price reductions (Table 2). The floral group rises in average price because the composition shifts from winter to summer dresses, that are more expensive in this example. This grouping is too broad; products not similar in price and purpose are grouped together causing the index to change because of composition rather than price.

Table 3. Grouping together dissimilar products can cause unintuitive price changes

Product	Price, Jan	Price, Aug	
Floral winter dress 1	18		
Floral winter dress 2	18		
Floral winter dress 3	24	18	
Floral summer dress 1	60		
Floral summer dress 2		45	
Floral summer dress 3		45	
Party midi dress 1	100		
Party midi dress 2		90	
			Price change (Aug/Jan)
Average price: Floral dresses	30	36	1.2
Average price: Party midi dresses	100	90	0.9
			Price index 1.05

Source: Office for National Statistics

5 . Assessing product grouping

The objective of product grouping is to create groups that are large enough to control for product churn but of a similar enough quality (known as "homogeneous") that compositional effects do not bias inflation. To measure these competing goals, [Chessa \(2019; PDF, 860KB\)](#) introduced "Match Adjusted R Squared" (MARS). MARS is the product of two components, both in the range [0, 1]:

$$MARS_{t0,t} = (match\ rate)_{t0,t} * (R - squared)_t$$

Where:

- Match rate measures product churn as the proportion of matching products or groups from the base to current month¹
- R-squared measures in-group price similarity within the current month

MARS accounts for homogeneity within groups based on price similarity, but this is only one measure of homogeneity. Guidance from the International Monetary Fund recommends [homogeneity should also account for purpose](#). For example, a group containing a £20 t-shirt and a £20 pair of shoes are homogeneous in price, but not purpose. We are exploring measuring purpose homogeneity through human evaluation of product similarity within groups, but this is ongoing research and is not presented in this article.

Notes for: Assessing product grouping

1. Chessa (2019) uses expenditure to weight the match rate. Since we lack this information in web scraped data, we use an unweighted variant.

6 . Attribute-based product grouping

Attribute-based product grouping creates groups of products that share characteristics. For clothing data, we primarily form attributes through text matching (for example, whether products contain the word "cotton"). A simplified example of this approach is provided in Table 4.

Table 4: Groups are formed from the attributes "polyester", "v-neck", "cotton" and "maxi"

Product name	Material	Group
v-neck dress	polyester	polyester_v-neck
floral maxi dress	100% cotton	maxi_cotton
floor length maxi dress	cotton, elastic	maxi_cotton

Source: Office for National Statistics

We need grouping models for over 100 clothing types such as women's dresses, boys' t-shirts and men's jeans. While models can be created from user-defined words, an automated approach is preferred for scalability.

For each clothing type, we generate groups using the most-commonly occurring words. Note that common words are likely to express quality-defining characteristics of products, for example, style, material and colour. To improve the grouping, common non-quality defining text are removed, including punctuation, numbers and stop words (such as "and", "but", "to"). We choose how many of the top words to use based on performance, currently assessed through MARS.

This method is relatively simple to deploy, scalable and we have found it has some desirable advantages relative to more advanced clustering methods that we have tried. More detail on this approach and our clustering approach can be found in [Dealing with product churn in web-scraped clothing data: product grouping methods \(PDF, 416KB\)](#).


7 . Results

We fitted the model using web-scraped data from June to December 2020, which determines the most common words and how many to use (with a longer data supply, we will use a minimum of a year to capture seasonal variation). We then applied the model to data from January 2021 to August 2021 to create product groups, measure "Match Adjusted R Squared" (MARS) and produce price indices.

In Figure 3, MARS scores are shown for women's dresses after applying product grouping. The method scores reasonably high on both the match rate and R-squared components. However, these results are experimental as we are continuing to refine our methods.

Figure 3: MARS scores for women's dress groups January 2021 to August 2021

Figure 3: MARS scores for women's dress groups January 2021 to August 2021



Source: Office for National Statistics

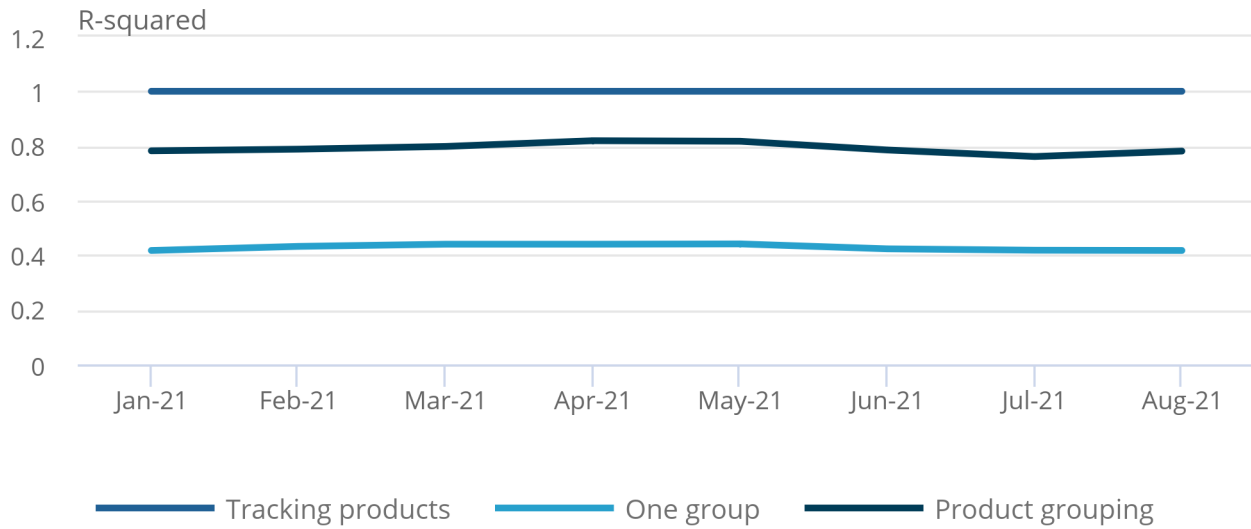
In Figures 4 and 5 we compare our R-squared and match rate against two benchmarks: tracking individual dresses and tracking a single group of all dresses. A perfect R-squared can be achieved by tracking individual dresses (Figure 4), but this results in a low match rate (Figure 5). Conversely, a perfect match rate can be achieved through grouping all products into a single group (Figure 5) but results in low price similarity within this single group (Figure 4). Our product grouping for dresses is shown to balance these two desired features of products used to construct a price index.

Figure 4: Attribute-based product grouping creates dress groups with greater price similarity than tracking a single group of all dresses

R-squared measures similarity of prices within groups

Figure 4: Attribute-based product grouping creates dress groups with greater price similarity than tracking a single group of all dresses

R-squared measures similarity of prices within groups



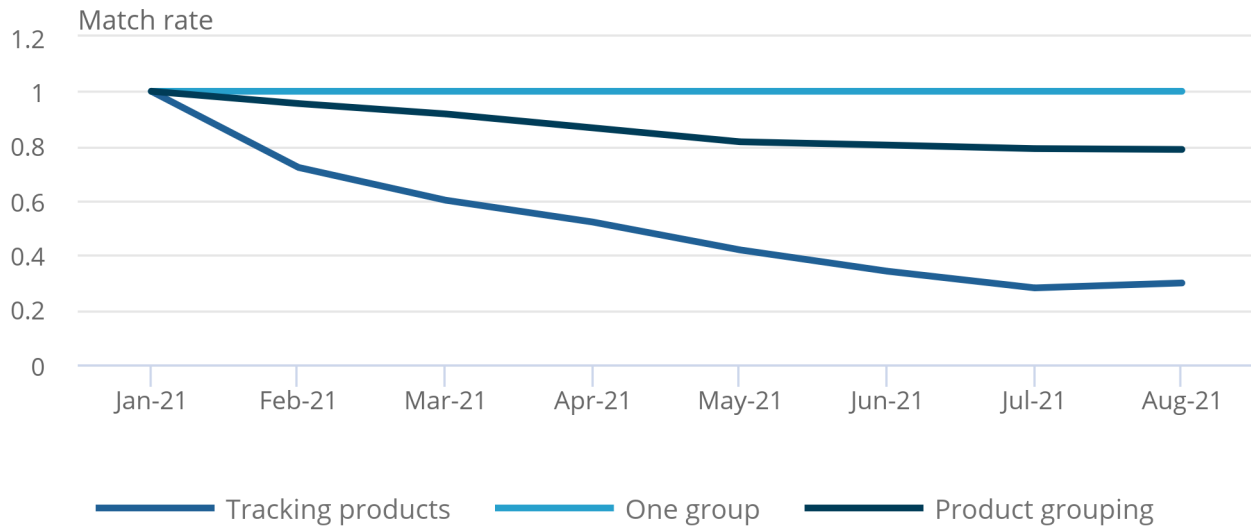
Source: Office for National Statistics

Figure 5: Attribute-based product grouping creates dress groups with a higher match rate than tracking individual dresses

Match rate measures how long groups remain on the market

Figure 5: Attribute-based product grouping creates dress groups with a higher match rate than tracking individual dresses

Match rate measures how long groups remain on the market

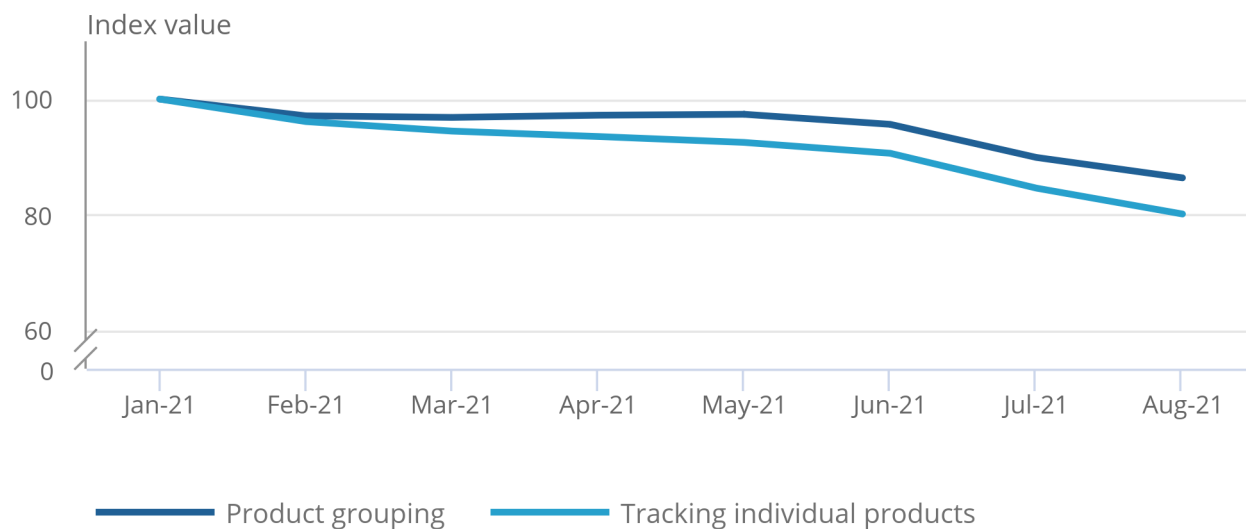


Source: Office for National Statistics

Figure 6 compares the price index before using product grouping (as shown in Figure 1) with the index after using product grouping. The method appears to capture the same seasonal patterns as not applying grouping but mitigates some of the fall in the index. Our future work looks to improve these methods and test them over longer time series of data.

Figure 6: Product grouping appears to capture the same seasonal patterns as tracking individual products but mitigates some of the fall in the index

Figure 6: Product grouping appears to capture the same seasonal patterns as tracking individual products but mitigates some of the fall in the index



Source: Office for National Statistics

8 . Future developments

In the coming months we will:

- use more data to both fit our model (to ensure variation in seasonal terminology is captured) and extend the index data time series to understand the longer-term impact of product grouping on indices
- productionise the system to ensure the method and these analyses can be scaled and applied to other clothing consumption segments
- make further refinements to the method with the goal of improving "Match Adjusted R Squared" (MARS); in particular refining which words are chosen to form groups
- explore measuring purpose homogeneity to determine whether the groups formed make an intuitive sense to consumers
- consider how we can better account for consumption patterns within the index, by applying retailer market shares and weighting individual products, or product groups, by estimates of consumer expenditure

There is still much to explore before clothing data can be included in our headline measures of consumer price statistics. Therefore, although we begin implementation of alternative data sources into consumer price statistics from 2023, the earliest that web-scraped clothing data will be included is 2024, after a period of sufficient development and testing.

9 . Related links

[Transformation of consumer price statistics: November 2021](#)

Article | Released 9 November 2021

Our plans to transform UK consumer price statistics by including new improved data sources and developing our methods and systems for production from 2023.

[Introducing alternative data into consumer price statistics: aggregation and weights](#)

Article | Released 9 November 2021

Plans to incorporate new data sources and methods into the existing structure of UK consumer price indices from 2023, including changes to the existing hierarchy and methods of weighting different strata.

[Consumer price inflation](#)

Bulletin | Released 20 October 2021

Price indices, percentage changes, and weights for the different measures of consumer price inflation.