

**Constructing variance estimates for the UK Consumer Price Indices:
Progress update April 2026**

Status: Work in progress

Expected publication: Alongside minutes

Purpose

1. The UK Statistics Authority, in approving the re-designation of CPIH as National Statistics, included a requirement to explore and publish estimates of quality. Since then, work to assess the sampling variance of the consumer price indices has been underway.
2. This paper summarises the progress made since July 2025.
3. In particular Annex A revisits the paradigm for variance estimation in the CPI, and explains the implications for the interpretation and explanation of variance estimates once they are available.

Actions

4. Members of the Panel are invited to:
 - note the progress with variance estimation
 - comment on the paradigms for variance estimation (Annex A)
 - comment on the treatment of prices with replication factors >1
 - consider how variance measures constructed under one or a combination of paradigms could be explained to both expert and lay users of the CPI.

Details

5. Since the previous update in April 2024 (APCP-T (25)09) we have:
 - worked on including the chainlinking process in the calculation of the variance of the CPI. This work is not yet complete.
 - Considered further the properties of the jackknife variance estimators of the elementary indices in design- and model-based situations; some additional thinking in this regard is presented in Annex A.
 - Begun consideration of the models which might be used for model-based variance estimation with scanner data (and similar data sources) and multilateral indices.
6. We continue to follow the plans for development of variance estimates. Some additional components are now in need of development. The following was highlighted in the previous paper, but remains an important step which is yet to be taken:
 - The work done by Jim O'Donoghue at the beginning of investigations of the variance of the CPI is now very out of date, and does not adequately represent the variability in the index due to the current basket and weights. The next step is to make this into a repeatable calculation which can be applied to the LCF data in any year, in order to

produce variances of the index accounting for the sampling variability in the weights. There are some challenges with how to deal with the effect of input-output balancing, which is a component of the weighting process, and how to deal with weights sources which are not derived from the LCF.

Paul A. Smith, University of Southampton
April 2026

Paradigm for calculation of the variance of the Consumer Price Index

Introduction

Considerable progress has been made with the calculation of a variance for the CPI, but this has not yet led to a final set of estimates. The main focus has been on the variance of the 12-month change in the index, since this is the key statistic on which users focus, though it is certainly possible to apply the methods to obtain variances for the index itself, or for changes over other periods.

Review

Smith (2021) reviewed the strategies and methods for variance estimation in Consumer Price Indices (CPIs). Most of the development has taken place in the US, where there was a gradual transition to a fully probabilistic design from December 1963, eventually being realised in 1978. This provided a firm basis for variance calculation using a number of different methods and estimators, including balanced half samples, jackknife and bootstrap approaches, and these have been extensively explored in the US CPI.

A further approach has been to consider the CPI as a model-based parameter; Kott (1984) suggested a fully model-based approach to the calculation of the CPI, but this has not generally been taken up – in particular it would imply a different approach to calculation of the index than that prescribed in the ILO Manual. Nevertheless, it is possible and practical to treat the variance of the estimators that are recommended under the model-based framework. Valliant & Miller (1989) and Valliant (1991) presented one approach, and Zhang (2010) considered model-based approaches leading to the Carli, Dutot and Jevons elementary indices.

Other authors have approached the CPI as if it is a design-based statistic, and then sought to produce an appropriate variance – see for example Biggeri & Giommi (1987), Dalén & Ohlsson (1995). Some design-based approaches have been focused on the variability due to sampling variation in the weights (eg Balk & Kersten 1986), which arise from surveys with more typical design-based properties than the price collection designs.

First, we consider the paradigm for construction of a CPI, and the implications for calculating a variance.

Procedures for calculation of a CPI

The CPI manual (ILO et al. 2020) specifies several stages in the production of the CPI. The first is the sampling procedure itself, where it advocates the use of probability sampling as a means of controlling for selection bias in the way that prices are gathered. But it does acknowledge that various parts of the design may not be able to be selected using a strict probability design, and that various alternative procedures are then possible.

The second stage of construction that interests us is the calculation of the lowest level indices. In most cases these are the so-called elementary indices, which use only the price quotes because there is no quantity information at this level with which to form weights. This is an implicitly model-based estimation, where we assume that the prices contribute equally to the calculation. But let us consider briefly what is being ignored in this process

- at one level we are ignoring the differential impact of different products within the calculation – for example in “large loaf, white, unsliced” we would ignore that the price(s) for branded loaves represent a larger share of transactions than the price for an unbranded loaf. This is

ignoring the usual weighting information, which is largely unavailable at this level (where it is available, we may use it).

- we are also ignoring the sampling weights. Even in the situation that the design is strictly probability based (as in the US), we *still* calculate the index at the lowest level giving the prices equal weights in the elementary index.

In fact this is not quite true, because an element of weighting is included in the calculations through replication factors (eg prices from store chains with national pricing are collected once, but given replicate weights to say how many times they should be used, reflecting the market share of the chain). We have so far assumed that replication factors >1 reflect that such stores are certain to be included (because of their market share), and therefore that they do not contribute to the variance.

Does the panel agree that prices with replication factors >1 should not contribute to the design-based variance?

Although these weights are ignored, there may still be a benefit in the probability design – ignoring the sampling weights may act as a way to give different price observations a different implicit contribution to the estimator, because of their differential selection probabilities. We expand on this a little further below.

The third stage of construction consists of a hierarchical aggregation of the elementary indices, and these are weighted aggregations, using information about expenditure proportions as the weights. Here we should account for the variability of the weights, which by default are generated under the design-based framework.

This calculation procedure defines the estimate of the Consumer Price Index. It completely eschews the explicit use of the sampling information for the prices, even though probability sampling is recommended as a basis for obtaining the price information. This creates a conundrum. We are used to the idea that an estimate from a survey is an estimate of a population quantity – a finite population parameter (FPP). And it is usually important to account for the differential selection of sample units through weighting in order to have an unbiased estimator of that FPP. But the CPI is not defined this way. There is no FPP that is our target parameter – and therefore no independent view of what constitutes “inflation”. Instead, the definition of inflation is given by the way in which the index (and therefore the change in the index) is calculated. In this sense the ILO manual (and the more strict control of the Harmonised Index of Consumer Prices in the European Union) are important for comparability, because they define the measure of inflation by the way in which the index is constructed.

Envisaged design

In the US CPI, all the stages of the design are probabilistic, so it is possible (at least in principle) to calculate the selection probability for any price quote included in the index. In the UK (and many other nations’) CPI, some of the stages are probabilistic, but some involve purposive selections (particularly the selection of representative items and the selection of products to price in the UK case. In these cases we do not have a selection probability, so need to approximate it in some fashion in order to understand the impact of the design on the variance of the estimator. Several authors propose the principle of pseudorandomisation, where the achieved prices are considered to be a random sample, often an equal probability random sample (ie a simple random sample) from the population. But we can also consider other situations which do not give rise to equal probability samples, and so we introduce the idea of an *envisaged design*, a design with which we assume the observations in the sample are generated. This allows us to evaluate the effect of the design assumptions on the variances, and to reflect realistic elements of the sample selection in the variance calculation where they are available.

In fact, we can see the careful design process as a way to implicitly include weights in the unweighted elementary index calculation. A price that has a higher sampling weight in the design will, in general,

be associated with a higher expenditure proportion (expenditure weight), and the realised sample will have more prices with higher expenditure weights since these also have higher selection probabilities. So, over repeated sampling, prices with higher expenditure appear in the sample more often. Of course this does not work very effectively in a single sample, but is still better than a completely equal probability approach. When using the unequal probability sample, it is important *not* to use the weights, so that the implicit weighting operates as expected.

Implications for variance estimation

There are at least two paradigms by which we can evaluate the variability within the CPI due to the price information.

First, we can view the whole procedure as model-based. If we set up a suitable model for which the calculated estimator is the outcome, or for which the outcome can be expressed in terms of the model parameters, we can evaluate the variability with respect to that model. This is the idea behind the model-based procedures and the estimation of model-based variances.

The second paradigm is to follow a design-based procedure, but with the (unweighted at the elementary level) calculated estimator as the target. In the usual manner we imagine a repeated sampling from the population using a fixed design. In each repeated sample we apply the procedure for calculating the CPI. We can then examine the distribution of this calculated CPI over the repeated samples, and the variability of this distribution is the variance that we wish to calculate – a design-based variance for a statistic which is not a FPP.

This dichotomy has not been explained in the literature. The model-based approach is reasonably internally consistent, matching the procedures used for estimation, and therefore seems a natural approach. It also creates a suitable structure within which other model-based adjustments can be handled (eg imputation, hedonic pricing, quality adjustment). The design-based approach is, however, rather peculiar in this framework. This has not stopped several authors from proposing procedures to deal with it (Biggeri & Giommi 1987, Dalén & Ohlsson 1995). Sitter produced exploratory variances for the UK RPI based on jackknife estimators, without ever commenting on the form of the estimation target that underlies it.

We are leaning towards describing all the variance estimation for the prices under the model-based framework, which also looks to be a practical approach for including some measure of variance for the index produced with scanner data.

Does the panel agree that a model-based approach is appropriate for the price variance of the CPI?
--

Should prices with replicate weights contribute to the variance under a model-based framework?
--

Accounting for the variance due to the weights

By contrast with the price collection and elementary index construction, which seem clearly closest to the model-based paradigm, the weights are derived from a survey in which the estimators are design-based, accounting fully for the sampling weights in the estimation process. There is some post-processing of estimates by balancing through the supply-use framework. It is not yet clear whether we can account for the effect of this process in variance estimation, which might be considered a model adjustment to the estimates from the survey. But initially it seems that a design-based approach to the variance of the weights is appropriate.

The variance of the weights is only part of the output however – we want the variance of the CPI with respect to both prices and weights, so the variances need to be combined. Analytically we would like

the variance with respect to the model for the prices and the design for the weights, but this can be more easily approximated numerically.

Interpreting a CPI variance estimate

These views of the way in which a variance estimate is constructed also affect how it can be communicated. The classical design-based approach would be an explanation something like “If we were able to repeat the sampling procedure many times, the sampling variance says how much the estimate would vary from the true population parameter”. But here there is no “true population parameter”, so at the least we would need to replace this by “...vary from the average population parameter over the repeated samples”. If we calculate a standard deviation for the 12-month change in the CPI, then how would we present this?

We could produce an alternative variance using a model-based approach, and this seems more in line with the construction methods for the index – even more so to evaluate the quality while using scanner data, as explored by the committee previously in ACPCP(25)09. Under the model-based scenario, the variance says nothing about repeated sampling. Instead it summarises how the observed prices differ from the postulated model. This naturally depends on the model chosen and how well it fits the data.

Some exploration of alternative models in order to find the form that best fits the data is in order. Ideally we would use a simple model that is flexible enough to capture the features in the data (for example correlations between the prices of different products in the same outlets, or correlation of prices over time).

How would such a model-based variance be explained to users? It captures the variability of the data relative to the chosen model and the CPI estimator. It doesn't have a repeated-sampling interpretation, but has a superpopulation rationale – if there were another UK in a parallel universe where the prices followed the same model with similar fit/variance, then the difference between the two would be reflected by the estimated variance. It is more difficult to explain/justify this approach to users.

What is the committee's view on how a variance estimate should be interpreted/presented?
--

References

- Balk, B.M. & Kersten, H.M.P. (1986) On the precision of consumer price indices caused by the sampling variability of budget surveys. *Journal of Economic & Social Measurement* **14** 19–35.
- Biggeri, L. & Giommi, A. (1987) On the accuracy and precision of the consumer price indices: methods and applications to evaluate the influence of the sampling of households. *Bulletin of the International Statistical Institute* **52** 137–154.
- Dalén, J. & Ohlsson, E. (1995) Variance estimation in the Swedish consumer price index. *Journal of Business & Economic Statistics* **13** 347–356.
- ILO, IMF, OECD, UNECE, Eurostat & The World Bank (2004) *Consumer Price Index Manual*. International Labour Office, Geneva. Available from <https://www.ilo.org/sites/default/files/wcmsp5/groups/public/%40dgreports/%40stat/documents/publication/wcms761444.pdf>
- Kott, P.S. (1984) A superpopulation theory approach to the design of price index estimators with small sampling biases. *Journal of Business & Economic Statistics* **2** 83–90.
- Smith, P.A. (2021) Estimating sampling errors in Consumer Price Indices. *International Statistical Review* **89** 481-504. doi: 10.1111/insr.12438.
- Valliant, R. (1991) Variance estimation for price indexes from a two-stage sample with rotating panels. *Journal of Business & Economic Statistics* **9** 409–422.

- Valliant, R. & Miller, S. (1989) A class of multiplicative estimators of Laspeyres price indexes. *Journal of Business & Economic Statistics* **7** 387–394.
- Zhang, L.-C. (2010) A model-based approach to variance estimation for fixed-weights and chained price indices. In *Official Statistics: Methodology and Applications in Honour of Daniel Thorburn*, Eds. M. Carlson, H. Nyquist & M. Villani, Stockholm University & Statistics Sweden: Stockholm, pp. 149–166.